Pervasive User Modeling and Personalization (PUMP 2010)

Workshop at

User Modeling, Adaptation and Personalization (UMAP 2010)

Big Island of Hawaii

June 20, 2010

Shlomo Berkovsky[1], Fabian Bohnert[2], Francesca Carmagnola[3], Doreen Cheng[4], Dominikus Heckmann[5], Tsvi Kuflik[6], Petteri Nurmi[7], and Kurt Partridge[8]

[1] CSIRO, Australia, [2] Monash University, Australia, [3] University of Torino, Italy, [4] Samsung Research, USA, [5] German Research Centre for Artificial Intelligence DFKI GmbH, Germany, [6] The University of Haifa, Israel, [7] Helsinki Institute for Information Technology HIIT, Finland, [8] PARC, USA

1 Introduction

Mobile and pervasive computing technologies have become an integral part of everyday life and have changed the way people interact with information. The rapidly growing amount of information and services raises the need for user modeling and personalization in pervasive and ubiquitous environments. Unique challenges include (1) inferring relevant information about a user from sensors, (2) aggregating and integrating such information effectively into long-term user models, while achieving user model interoperability, and (3) providing pervasive and ubiquitous information access in a personalized manner. The objective of this workshop is to bring together active researchers and practitioners working on user modeling and personalization in pervasive and ubiquitous environments and to produce vision statements about the future of these fields.

Workshop Topics. The workshop addresses the following four focus questions:

1. What pervasive or situational information is most useful for user modeling and personalization, how can such information be extracted from sensors, and how can the information be represented and effectively used in user models?
2. How can we aggregate and integrate user modeling data from various sources, resolve conflicts, and abstract from and reason about the data?
3. How can we overcome syntactic and semantic heterogeneity of distributed user models in order to achieve user model interoperability?
4. What unique challenges and opportunities exist in pervasive ubiquitous environments that are not present in conventional online personalization domains, and how can these challenges be addressed?

1.1 Organizers

- Shlomo Berkovsky, CSIRO, Australia
- Fabian Bohnert, Monash University, Australia
- Francesca Carmagnola, University of Torino, Italy
- Doreen Cheng, Samsung Research, USA
- Dominikus Heckmann, German Research Centre for Artificial Intelligence DFKI GmbH, Germany
- Tsvi Kuflik, The University of Haifa, Israel
- Petteri Nurmi, Helsinki Institute for Information Technology HIIT, Finland
- Kurt Partridge, PARC, USA

1.2 Program Committee

- Lora Aroyo, Free University of Amsterdam, The Netherlands
- Jörg Baus, German Research Centre for Artificial Intelligence DFKI GmbH, Germany
- Federica Cena, University of Torino, Italy
- Nadja Decarolis, University of Bari, Italy
- Eyal Dim, The University of Haifa, Israel
- Judy Kay, University of Sydney, Australia
- Gerrit Kahl, German Research Centre for Artificial Intelligence DFKI GmbH, Germany
- Matthias Loskyll, German Research Centre for Artificial Intelligence DFKI GmbH, Germany
- Antti Oulasvirta, Helsinki Institute for Information Technology HIIT, Finland
- Francesco Ricci, Free University of Bozen-Bolzano, Italy
- Rainer Wasinger, University of Sydney, Australia

Table of Contents

ARBUD: An Architecture for Building Pervasive User Models from Massive Sensor Datasets

Heath Hohwald, Enrique Frias-Martinez, and Nuria Oliver

Data Mining and User Modeling Group
Telefonica Research, Madrid, Spain
{heath,efm,nuriao}@tid.es

Abstract. In pervasive environments it is common that data from a large number of heterogeneous sensors serves as input for generating a large number of user models for applications with time constraints. This situation raises the need for an architecture that can be employed for efficiently constructing the variety of user models needed by different applications. In this paper, we propose a distributed-computing architecture based on MapReduce that allows for the efficient processing of massive and heterogeneous sensor datasets using reusable components. A metamodel is used for specifying the characteristics of the desired pervasive user model, which can include both short-term and long-term features, and a library of reusable components factors out commonality across applications and sensors. We present an instantiation of the architecture for generating user models for mobile phone subscribers and empirically evaluate the scalability of the proposed architecture with a large real dataset. Our results indicate that complex pervasive user models for millions of users and thousands of sensors can be obtained in just a few hours on a small computer cluster.

1 Introduction

One of the areas in which pervasive computing has had a significant impact is in public environments, where a wide variety of information is transparently collected from users. Information collected includes public transportion, mobile phone data (both voice and web navigation), traffic data, etc. The information available for each user originates from different sensors that produce large quantities of information about individual behavior, raising the need for platforms and architectures that are able to efficiently generate pervasive user models (UM) for a large number of users from massive and heterogeneous datasets.

Until recently, the user modeling literature has not devoted much attention to scalable architectures for efficient and large-scale pervasive user modeling, mainly because most approaches assume that all the information to be processed is located at a specific, limited device, e.g. a mobile phone. Such an approach has two major limitations: (1) applications and models are highly dependent on the platform and its capabilities and (2) the models are available only locally, while aggregating such information is useful for applications relating to smart cities and smart environments. Also, more scalable approaches have not been considered because the number of sensors and the number of users available have been limited. Nevertheless, the increasing availability of large scale human activity data in a variety of pervasive domains (*e.g.,* the mobile Web, telecommunications, traffic, public transportation, etc.) creates the need for platforms that can centrally and efficiently generate large numbers of complex user models.

The work most related to our research is in the field of data mining and web mining frameworks [8], where some previous efforts have also focused on high performance [5]. However, these frameworks are not generally geared towards user modeling. In this context, [7] presents WIM (Web Information Mining) for web mining prototyping and provides a model and an algebra to express web mining tasks.

In this paper we present an architecture that centrally generates user models from the datasets produced by a variety of sensors of a pervasive computing infrastructure. Based on the MapReduce paradigm [4], the architecture: (1) is able to efficiently generate complex user models in a timely fashion from massive amounts of sensor data and for a large numbers of users; (2) has the capability of generating both short and long term user models, and (3) uses pluggable components where the functions that are used to compute each of the features in the user model can be tailored to the different sensors, can be reused from existing libraries, and can be shared across multiple application domains. Note that our approach is very different from updating a user model on-the-fly when new information arrives, as done in AHA! [3]. Such an approach requires a much more complex data representation and has to be built using problem-specific architectures. In general, these solutions are not designed for the characteristics of pervasive environments. To the best of our knowledge, the proposed architecture is the first of its kind to address the challenges of scalability, reusability and domain independence for constructing user models from large heterogeneous sensor data generated from pervasive environments.

2 Architecture for Terabyte-Scale Pervasive UM

The proposed architecture for building pervasive user models, named ARBUD, is depicted in Figure 1 and has four main components: (1) The *Sensor Data* available to generate user models, which can be one or more sources of data; (2) a *User Metamodel* that describes the components of the long term (LT) and short term (ST) user models; (3) a *user modeling (UM) Library* that contains the most common functions used to generate the different features of the user models, and (4) a *UM Generator* that instantiates the architecture and generates the set of pervasive user models.

Data Input & Output. The sensor data available to generate the user models typically is available as a set of files, with one or more files for each sensor. The output of the architecture can be presented and managed as a single file, but for scalability purposes typically it will consist of a set of distributed files along

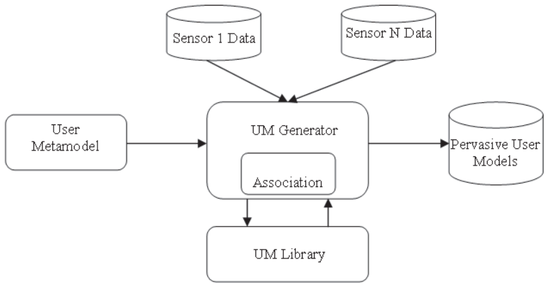

Fig. 1: Proposed architecture.

with a lookup table that indicates which file contains the associated user model for each user.

User Metamodel. The metamodel contains all the parameters that describe the desired pervasive user model at a high level. This metamodel allows the architecture to factor out features common to several application areas or sensor sources. It also allows for the easy use of pluggable components tailored to the different types of sensors. As a consequence, models can be built without detailed knowledge of the data layout or the distributed programming code underlying the architecture. The metamodel contains the following parameters:

1. *INPUT* is a pointer to the directory or files that contain the sensor data to be processed. Typically, each file comes from a single sensor and will contain the activity logs for all users during a given period of time. Each entry in the file will have a reference to the user ID and the user's activity in a specific moment in time.
2. LT_s, LT_e, ST_s, and ST_e correspond to the start and end times used to generate the Long Term (LT) and the Short Term (ST) models. Setting $LT_s = LT_e$ only generates a ST model and likewise for LT.
3. LT-UM is a vector ($\text{LT-UM}_1, \cdots, \text{LT-UM}_N$) indicating the features that define the long term user model. The processes for generating these features in a given pervasive environment can be mostly standardized and are included in the UM Library.
4. ST-UM is a vector ($\text{ST-UM}_1, \cdots, \text{ST-UM}_M$) indicating the parameters that define the short term user model. Note that LT-UM and ST-UM can be different and can have a different number of features. Often they share many dimensions, but the short term model may include dimensions that only have significance over a short term period.

UM Library. The UM Library contains the functions needed to calculate the features specified in the LT-UM and ST-UM. Note that for efficiency purposes, one or more features may be the output of a single function. For a given context, the UM Library will contain the typical elements that user models include in that environment. For example, for data originating in mobile web navigation there can be a function to identify the number of times a user visits a web page according to the user's location. For data originating from mobile phone towers, a function might calculate the total talk time for a given subscriber in each location or the number of individuals in a subscriber's social network.

UM Generator. The UM Generator applies the User Metamodel to the data originating from the set of sensors and produces the Pervasive User Models. The Generator matches each of the elements of ST-UM and LT-UM to the corresponding functions of the UM Library with a lookup table; for example, in the case of LT-UM it maps each dimension LT-UM$_i$ to a function F_i. Figure 2a depicts simplified pseudocode for the UM Generator, which builds the short term and long term user models in parallel (the outer **ParFor**s). Within each user model, all features are computed in parallel (the inner **ParFor**s). The Generator waits until all features are computed, and then aggregates the results.

Figure 2b presents an example of the workflow involved in building a LT-UM. The metamodel specifies the location of the input data, typically a set of time-stamped sensor log files. The architecture then launches in parallel all functions $F_1, F_2, \cdots, F_N$ each of which produces a set of files containing records that contain a user ID and the associated feature. The architecture waits for all functions F_i to complete and then launches another process to aggregate all the results. The aggregation process is responsible for merging all of the values for each of the i dimensions associated with each user. It produces as output a series of final LT-UM files where each record consists of a user ID and the full vector of N associated features.

```
ParFor
        ParFor i=1...M
                Launch(ST-UMi)
        EndFor
        Wait(ST-UMi) i=1...M
        Aggegate (ST-UMi)

        ParFor i=1...N
                Launch(LT-UMi)
        EndFor
        Wait(LT-UMi) i=1...N
        Aggegate (LT-UMi)
EnfFor
```

(a) Pseudocode

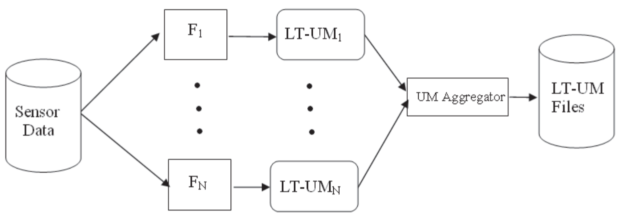

(b) Example Workflow for Generating LT-UM

Fig. 2: Pseudocode and Example Workflow for the UM Generator

Implementation. The proposed architecture is heavily influenced by the concepts and philosophy of the MapReduce [4] programming model and the accompanying distributed file system. MapReduce is a framework for processing massive amounts of data in parallel using a (typically large) cluster of computers. The data can be unstructured which is ideal for a pervasive user modeling applications. MapReduce works with key-value pairs and consists of two basic stages: (1) A mapping stage that processes input records in parallel and emits intermediate key-value pairs, and (2) a reducing stage that gathers all values associated with a given key and then reduces the set of associated values, emitting final key-value pairs. For the architecture considered here, the keys are often user IDs and the final values are typically one of the features in the user model, so MapReduce is a natural choice for constructing the UM Generator. We have implemented the UM Generator using Hadoop [1], which has two core components: an open-source Java implementation of MapReduce and HDFS, the associated distributed file system. The functions included in the UM Library will usually be implemented as MapReduce processes, although they can be implemented using traditional programming paradigms.

3 ARBUD Applied to Mobile Phone Sensor Data

In order to evaluate the performance of the proposed architecture and to demonstrate the advantages of having a reusable UM library, we illustrate how ARBUD can be used to centrally build user models from a pervasive environment. In this case, the sensors are mobile phone towers, which are ubiquitously located and collect a large amount of information about each individual including location, phone call behavior, social network and mobile web navigation features. The information captured by the sensors is stored in data files that the serve as the input for our example. In general, a large number of applications of user modeling in telecommunications, such as churn or location services, focus on understanding how individuals use their phone, and how that knowledge can be used for providing better services. In these applications it is common to process very large datasets with tens of millions of users and where obtaining user models as quickly as possible is critical. For example, churn or fraud detection models are typically run on a daily basis in order to detect users at risk. Also, the different user models for each application can have elements in common which can be shared using the reusable nature of the proposed architecture.

Mobile Phone UM Library. The Mobile Phone UM Library assumes that each mobile tower records its own set of files, though additional logs coming from other sources, such as mobile transactional payment logs, would be handled in the same fashion. The UM Library we have created for this domain includes numerous functions. Among others, there are functions to calculate, for each user, (1) the total number of calls made, (2) the total number of calls received, (3) the total duration of the calls made, (4) the total duration of the calls received, (5) the set of locations preferred by the user, (6) the frequency with which the user changes terminals, (7) the in, out, and total degree of each user, where degree captures the size of a user's social circle, (8) the number of calls to other service providers, (9) the number of calls to customer service, (10) number of individuals in the user's social network that use another carrier, (11) the number and duration of calls made for each hour of the day, and (12) the number and duration of calls made for each day of the week.

Mobile Phone Metamodel. Using the Mobile Phone UM Library, a metamodel can be constructed and used to generate long and short term user models for different applications. For example, churn [2] is one of the main problems of any telecommunications operator and it is defined as the percentage of customers that leave the operator in a pre-defined time period. In general, the reasons for churn are varied, but there are a number of factors that are good indicators of churn, such as the number of individuals in a subscriber's social network that use another carrier, the number of complaints the subscriber makes to the operator, or a large number of phone calls to other service providers. The user model generated by our architecture will be the input to an SVM that is trained to identify people at risk of churning. The classification needs to happen on a daily basis, thus the need to have models of all users in a timely fashion (*i.e.,* daily in this case). In this context using a ST Model and a LT Model is very useful for measuring recent changes in individual behavior and we assign ST_s and ST_e to span 5 days and LT_s and LT_e to span 30 days. ST-UM and LT-UM in this case are identical, and consist of the set of functions $\{1, 2, 3, 4, 6, 7, 8, 9, 10\}$ from the Mobile Phone UM Library.

4 Performance Evaluation

In order to evaluate the performance and scalability of ARBUD, a reference implementation was developed using a combination of Java and Hadoop. The performance of the architecture, measured in terms of total running time (in seconds) to produce user models, was evaluated with respect to: (1) the number of CPUs available in the compute cluster, and (2) the number of dimensions in the short and long term user models. For the purpose of the performance evaluation, anonymized phone call detailed records (CDRs) originating from the users and collected by the mobile phone towers (sensors) were used.

From the Mobile Phone UM Library described above, seven variables $\boldsymbol{v} = (v_1, v_2, \cdots, v_7)$ were chosen that exhibited similar performance characteristics. Using the randomly chosen set of variables $\boldsymbol{v}$, 14 metamodels were created, 7 consisting only of a short term user model and 7 consisting only of a long term user model. The short term metamodels, denoted as $MM_S^1, MM_S^2, \cdots, MM_S^7$, specified 3 days of data comprising $2.18 * 10^8$ records collected by the sensors with the ith short term metamodel MM_S^i denoting a user model consisting of the first i variables $(v_1, v_2, \cdots, v_i)$. The long term metamodels, denoted as $MM_L^1, MM_L^2, \cdots, MM_L^7$, were built from 30 days of sensor data, $2.17{*}10^9$ records and 103.0 GB. Similar to MM_S^i, long term metamodel MM_L^i included the first i variables of $\boldsymbol{v}$. The data was generated by 20 million users and their activity was collected by 20,000 sensors. The experiments were run on a compute cluster with five nodes. Each node consisted of 16 GB of RAM, 4 hard drives each with 1 Terabyte storage capacity, and 4 quad core processors.

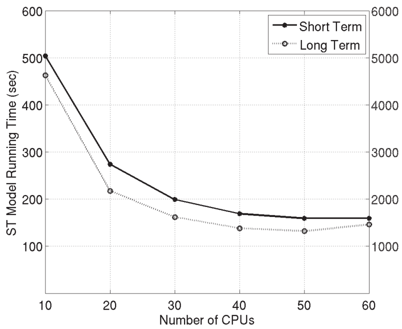

(a) Time vs. Number of CPUs

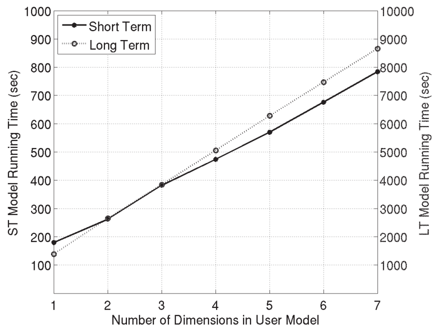

(b) Time vs. User Model Size (40 CPUs)

Fig. 3: Time needed to build long and short term user models.

Performance as a Function of Number of CPUs. To test scalability of the architecture for generating pervasive UM, different configurations were used in order to disable different numbers of the cluster's CPUs, though always ensuring equal CPU counts for each of the 5 machines in the cluster. For each configuration, one-variable long and short-term metamodels MM_S^1 and MM_L^1 were built using different numbers of CPUs. The results are presented in Figure 3a, where the x-axis reflects the number of CPUs enabled in the cluster, and the left (short term) and right (long term) y-axes show the total time taken in seconds to build MM_S^1 and MM_L^1. There is a factor of 10 difference in the left and right y-axes, reflecting the fact the the long term model is built from roughly 10 times as much data. In both cases, user model construction time initially decreases as more CPUs are added, then stabilizes with performance reaching a maximum at about 50 CPUs. This pseudo-exponential performance curve is in agreement with previous findings [6]. Even with only 10 available CPUs, more than 2 billion sensor inputs were processed and MM_L^1 was built in about 1.27 hours.

Performance as a Function of User Model Size. The performance of the proposed architecture scales well as additional features are added to the user model. To test the relationship between total running time and the number of dimensions in the user model, the cluster was configured to always use 8 CPUs for each of the 5 machines. The same 7 short and long term metamodels ($MM_S^1, MM_S^2, \cdots, MM_S^7$ and $MM_L^1, MM_L^2, \cdots, MM_L^7$) were each passed into the architecture and the corresponding user models were built. The resulting running times are presented in Figure 3b, where the x-axis represents the number of dimensions in the user model and the left (short term) and right (long term) y-axes show the total time taken in seconds to build the short and long user models. Variance was found to be minimal and results are shown from one set of models built. Both short and long term models are seen to exhibit linear scalability. As in Figure 3a, there is a factor of 10 difference in the left and right y-axes. The correlation coefficient between the number of model dimensions and the short and long term running times are 0.9994 and 1.0000, respectively, providing strong evidence for linear scalability. Building a long term model with 7 different features from a real data set with more than 2 billion records took a little more than two hours, implying that the architecture can easily rebuild the pervasive user models on a daily basis for the data sets considered. To further test scalability, a 7-dimensional user model from a 1.02 terabyte sensor dataset generated by 20 million users was built using ARBUD, employing 24.25 hours of compute time on a small 5 node computer cluster.

5 Conclusions

Pervasive environments produce a massive amount of data for potentially millions of individuals from a variety of heterogeneous sensors, highlighting the need for an architectures capable of processing all the sensor data and efficiently generating user models. In this paper, we have presented ARBUD, an architecture for building pervasive user models from massive sensor datasets that scales efficiently. The architecture uses metamodels to describe the desired user models, allowing for the abstraction of commonality from different user model building processes and different sensor functionalities. We have experimentally verified that ARBUD's performance scales linearly with the number of features in the user model and verified that ARBUD is able to process a terabyte of sensor data in 24.25 hours on a small cluster. These results highlight the ability of the proposed architecture to take advantage of recent advances in distributed computing to efficiently produce user models from large sensor datasets for pervasive applications with time constraints.

References

1. Hadoop information. `http://hadoop.apache.org`.
2. K. Dasgupta, R. Singh, B. Viswanathan, S. Mukherjea, A. Nanavati, and A. Joshi. Social ties and their relevance to churn in mobile telecom networks. In *Proc. 11th Int. Conf. Extending database technology*, pages 668–677. ACM, 2008.
3. P. De Bra, A. Aerts, B. Berden, B. De Lange, B. Rousseau, T. Santic, D. Smits, and N. Stash. AHA! The adaptive hypermedia architecture. In *Proceedings of the fourteenth ACM conference on Hypertext and hypermedia*, page 84. ACM, 2003.
4. J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. pages 137–150, 2004.
5. R. Grossman and Y. Gu. Data mining using high performance data clouds: Experimental studies using sector and sphere. In *Proceeding of the 14th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 920–927. ACM, 2008.
6. U. Kang, C. Tsourakakis, and C. Faloutsos. Pegasus: A peta-scale graph mining system - implementation and observations. *IEEE Int. Conf. Data Mining*, 2009.
7. Á. Pereira, R. Baeza-Yates, N. Ziviani, and J. Bisbal. A model for fast web mining prototyping. In *Proceedings of the Second ACM International Conference on Web Search and Data Mining*, pages 114–123. ACM, 2009.
8. C. Wood and T. Ow. WEBVIEW: an SQL extension for joining corporate data to data derived from the web. *Commun. of ACM*, 2005.

Lifelong Personalized Museum Experiences

Tsvi Kuflik[1], Judy Kay[2]and Bob Kummerfeld[2]

[1]University of Haifa, Israel
(tsvikak@is.haifa.ac.il)
[2]School of Information Technologies, University of Sydney, Australia.
{judy.kay,bob.kummerfeld@sydney.edu.au}

Abstract. Previous research on the personalized museum experience has largely focused on the single visit. New and emerging mobile technologies are enablers for a longer term view, where the personalization spans multiple visits to museums and links them to other aspects of the user's life. There is a large body of research about the importance of this form of personalization. This paper reviews key literature that informs the directions that seem most promising for "lifelong" personalization of the museum experience, identifies some of the technical challenges for such personalization, in terms of the user modeling, ontologies, infrastructure and generation of personalized content. By taking this approach, technology can play a major role in supporting users in their ongoing museum experience, by modeling the visitors, "remembering" their history and recommending a plan for future visits. We explore the role of client-side personalization that makes use of a mobile phone as a lifelong user model server and a platform for both interacting with a technology rich museum environment and for delivery of personalized information.

Introduction

This paper explores a little studied aspect of research into personalizing the museum experience. Essentially, we rethink the view of the personalized museum experience, to see a visit to a particular museum, at a particular time, as just one link in a "lifelong" chain of visits to museums, within one's life experiences. In some cases, the user will repeatedly visit the same museum. In other cases, particularly in the case of tourists, there may only ever be a single visit that museum. However, even in this case, it is valuable to see it as part of their total, long term museum experience.

To see why this seems an important issue, we first consider the substantial literature about the reasons that people visit museums. Much work refers to the museum visit as a form of *leisure* activity, for example in Falk and Dierking [1992]. People come to a museum with children, for *recreation*, because of the *reputation* of the museum or a current exhibition, out of *interest* (in the exhibition/museum/ museology) [Durbin, 1996] or simply because it is raining [Falk, 2009]. Falk, based on a lifetime of research in the area, concluded that the notion of *identity* is important for understanding the reasons that bring a person to a museum and defining their goals. He identifies five major categories of identity-related groups: Explorer, Facilitator, Experience seeker, Professional/Hobbyist and Recharger. He emphasizes that whatever the visitor sees or does is influenced by the combination of their identity, their personal context (prior knowledge, experience and interest), the physical context (the specifics of the exhibition he/she encounters) and the socio-cultural context (the within- and between-group interaction that occurs in the museum). Hence, the visitor perceives the museum experience as satisfying if there is a good match between the visitor's identity-related needs and the museum affordances. During and just after the visit, the visitor constructs meaning from the experience, where this is particularly related to identity building. If a system is to improve the visitor's experience, it should be able to do this better if it takes account of the visitor type, their personal context (prior knowledge, experience and interests), social context and the museum physical context.

Whatever reason brings a person to the museum, they have a limited amount of time. This is defined by various constraints, including, the opening hours of the museum, the time the person has available, their own attention span and that of their companions. It is well recognized that museum visitors typically experience “museum fatigue” [Bitgood, 2009]. This is an important factor for our “lifelong” view of the museum experience because each particular visit is time-limited. This makes it particularly valuable to help the visitor make the best of each visit and to achieve what they want from it, where this may change according to their current needs and companions as well as in light of their previous experiences. It seems particularly promising to explore how to do this for repeat visits to the same museum, and to museums that are quite similar or ones that are located close to each other and have related exhibitions.

The core of this paper will explore the ways that this understanding of the nature of the museum experience relates to the ways that previous research has tackled various elements of the personalization of the museum experience (even though the museum visit is often a social event, that calls for supporting groups rather than individuals, we focus at the moment on the individual visitor). In taking a “lifelong” perspective, we have identified four ways to view a museum visit, beyond the single visit view that has been the focus of most work on personalization of the museum experience:

- Single museum, repeat visits: where the user, perhaps with companions, comes back to the same museum, either to revisit the things that they particularly enjoyed last time or to see new things;
- Related museums: where there is potential to link the experience at this museum with previous or future visits at related exhibitions, for example groups of museums that are near each other;
- Independent museums: which are particularly important for tourists who may only be able to visit a particular museum once, but they may gain more from that experience if it is linked to their other museum experiences;
- Links to rest of user's life: where other aspects of the user's life are related to the personalized experience at this museum, a common and important case being the museum visits by school children where the links with classroom experience may be critical to the effectiveness of learning.

There are important differences between these views in terms of the need for interoperability as well as management of the user model. The list is ordered in increasing complexity, as will be discussed and explained later. The later views introduce the potential for new ways to link a visit to the rest of the user's life.

The next section provides a brief overview of related work. Then, we analyze the potential of emerging pervasive and ubiquitous technologies, combined with lifelong user modeling to underpin a new approach to creating a personalized museum experience. Then we identify key technical challenges and conclude with a summary.

Related work

There is a considerable research into personalization of the museum experience. Broadly speaking, each project has tended to focus on a particular technological challenge. For example, one important group of early work explored the generation of natural language for the delivery of personalized information (for example, ILEX [Oberlander et al, 1998] or Alfresco [Carini et al., 1993]). More recently context awareness, particularly exploiting location modeling, was explored for museum visitors guides [Not et al., 1998; Opperman and Specht, 1999]. Another important focus has been on the generation of multi-media presentations (such as in the PEACH project [Stock et al., 2007]). There has also been some work on supporting groups of visitors (for example, Pil project, [Kuflik et al, 2007]).

While most work has dealt with the museum visit as an isolated event, some work has considered the time before and after it. Visitors may want to keep information to help maintain memories of a visit. Several systems have helped visitors to compose a summary during the visit, to take/send home as a souvenir and to consult later, perhaps to delve further into the subject [Grasso et al., 1998; Garzotto et al., 2003]. PEACH [Stock et al., 2007] took this visit summary one step further, with the automatic generation of a personalized museum visit summary report. For the period before the visit, CHIP [Wang et al., 2009] supported a pre-visit website, where the users could plan their visit and, at the same time construct a user model that may be used later on to support on site tour.

A completely different aspect of research for this paper comes from the emerging possibilities for lifelong user modeling [Kay, 2008, Kay and Kummerfeld, 2009]. A primary reason for this is due to the technological progress and potential of pervasive computing. So, for example, it is becoming clear that people will carry a powerful personal computing device. This will be capable of sensing aspects of the context, such as the user's location and activity. It will be able to store a user model, communicate with the environment, for example to download personalized museum tours, and to perform client-side personalization [Kuflik and Poteriaykina, 2009; Gerber et al, 2010]. It will have the potential to store episodic user models, with links to rich materials that evoke them [Kay 2009]. This links well with current understanding of role of memories associated with museum visits [Falk, 2009].

Pervasive technologies for personalized museum experiences

There are several important classes of pervasive technology that will have an important impact on the ways that the museum experience can be enhanced with personalization. Falk [2009] emphasizes that the museum visit experience involves several personal, physical and social contexts, which interact with he visitor’s identity. The physical context is given. It is the environment we are in, but if monitored then the dynamic aspects of the physical context can be used to enhance the visit experience. This may include lighting conditions, how crowded and noisy the environment is. This may be combined with personal aspects, such as how sensitive the visitors are to noise and whether the presence of others disturbs them, for guiding the system in supporting the visitors. This form of personal information may be stored in the visitors’ lifelong user models, ready for the system to use. The personal context is where the lifelong user model can provide most of the information – the visitor's prior knowledge of the domain, interests and current preferences all may be represented, and updated in the user models, as well as their visit history to this and other museums, that may be used to create more coherent information and support during the visit. The social context itself may be supported by knowing the individual group members and their current social relationships and use this information to foster interaction for the current museum experience.

The emergence of “Social Signal Processing” [Vinciarelli et al., 2009], that is based on measuring and reasoning about “thin slices” of information in ubiquitous computing offers the potential for better monitoring of users in computerized environments and hence opens a wide range of possibilities for better support based on better understanding of visitors behavior, as suggested by Dim and Kuflik [2010].

We envisage that users will, increasingly, carry quite powerful personal devices, in the form of mobile phones. These can be expected to provide computational power, memory and a connection to the network that will play an important role for the lifelong museum experience. They are adequate for the delivery of multimedia information about the museum and they can support interaction with companions. They will have the power to store a user model [Gerber et al., 2010; Kuflik and Poteriaykina, 2009] and this can provide the foundation for client-side personalization. Smart phones can provide reliable outdoor location, but typically the resolution is no better than 10 meters. Even this is not available indoors and, even so, would be inadequate for museums in terms of granularity. Current phones are limited in indoor environments but there is considerable work to improve indoor localization.

Another emerging technology for enhancing the museum experience is embedded “surface” computing, as interactive tabletops and large wall displays. The user's smart phone can serve an important role, especially for sending data files, and possibly also parts of the user model to support surface computing interaction.

The interaction between the visitor’s profile (probably on the mobile device) and the environment with its embedded sensors and computers may be greatly enriched by being aware to the cumulative museum visit experience the visitor collected along his lifelong museum experience.

Technical challenges for personalized lifelong museum experience

There are key technical challenges for "lifelong" user models. One is the representation of the user. We need to define the user model ontology and the mechanisms for acquiring information about the user and reasoning about them. There is considerable research on user model representations and reasoning. To support the personalized museum experience, we need to draw upon approaches that can deal with the interaction between the current context and long term, more stable parts of the model, such as knowledge and interests. To populate the individual's model, the main approaches are:

- implicit, based on the user interaction with the personalized system, as in even the earliest systems such as Alfresco [Carini et al, 1993] or the overlay user model of PEACH [Stock et al., 2007] and CHIP [Wang et al., 2009].
- implicit, based on observations of the location of the user, for example, using vision technologies such as ec(h)o [Hatala and Wakary, 2005];
- explicit modeling based on a questionnaire, as in HyperAudio [Sarini and Strapparava, 1998] and the bootstrapping phase of CHIP [Wang et al., 2008];
- stereotypes, which may be selected by the user, as in PEACH [Stock et al., 2007] or by the system as in LISTEN [Zimmermann et al., 2005]
- models of what the system has told the user, so that it can avoid repetition as in M-PIRO [Androutsopoulos et al., 2007] but which might equally have a role in helping a user re-find things that they found interesting but where they have forgotten some of the details.

The key change that a lifelong user model brings is that the user can carry their user model, for example on their mobile phone [Gerber et al, 2010; Kuflik and Poteriaykina, 2009]. This reduces the need for acquisition of the model on each visit. So the model can then be reused in and from other contexts. The attractive possibility arises for a system to help the user see links between information presented in the museum and in other contexts, particularly if the user model includes an episodic model, which represents the vivid memories that people have of key parts of museum visits [Falk, 2009].

To model the user's context, particularly their location and relevant aspects of their activity and attention, there are significant technological challenges, both in collecting relevant information from sensors and then interpreting it. This may run on the user's carried device and/or an infrastructure of sensors in an instrumented museum, where visitors can be monitored, tracked and modeled. These are challenging preconditions for any museum: adding instrumentation may be unacceptable.

There are other ontological challenges. It may be possible to define agreed ways to represent the user model, such as in userML [Heckmann and Kruger, 2003], so that its syntax can be understood by independent parts of the infrastructure. Then there may be the need for ontology matching, if the museum and user model ontologies differ. Here interoperability of user models becomes a practical challenge – every museum may require specific information, in its' unique representation, that may differ in every case and from the personal lifelong model as well. Hence some form of user models mediation [Berkovsky et al., 2008] may be required as well. There is a need for effective infrastructure for managing distributed models across the museum's pervasive computing environment, as in Personis [Assad et al, 2007]).

Presentation generation is another major challenge, especially as people expect high quality. Previous work to create presentations that are coherent across the visit used sophisticated natural language generation that appears at the moment too complex for practical deployment. The task is much harder when we consider multi museum visits and lifelong experience, where presentations in one museum should take into account past experience in and outside museums and, possibly, suggest future experience elsewhere.

These requirements are indeed challenging. Table 1 shows the increasing complexity of different forms of longer term personalization. The key challenge that differs across the rows is the nature of the ontology that will be the basis of the domain, and hence, user model. Consider the first row, for repeated visits to a single museum: For all four aspects, represented by the columns, supporting it might be feasible with minor extensions to existing work, for example the approach taken by CHIP [Wang et al., 2009]. The key innovations are managing the long term model (perhaps on the user's phone) and maintaining a history of prior visits, then exploiting these for the information delivered.

Table 1. Museum Scenarios and Technical challenges

	User Model	Domain Ontology	Infra-structure	Content composition
Single museum, repeat visits	easy	easy	easy	easy
Related museums	less easy	less easy	Less easy	less easy
Unrelated museums	harder	harder	Less easy	harder
Links to the rest of user's life	very hard	very hard	harder	harder

For the second row, related museums, all aspects become somewhat harder. The key new problems relate to potential differences in domain ontology, but being topically related, these should be "close" and potentially amenable to automated mapping. Turning now to unrelated museums (e.g. Art versus Archeology or Science museums), in this case, the challenge is to make useful links even though the ontologies may be very different. So, for example, the subject of a painting may be related to displays in a science museum. Content preparation becomes much harder, because of the need to map the museum experience between different domains and scenarios. For the final row, linking to everyday life, things are more complex again. It is not clear how personal information should be stored and maintained for use by multiple systems over time. Perhaps at this level, the priority should be to support the user in re-finding things that they partly recall. A form of associative search on the user model might be helpful [Collins et al, 2009]. With technological changes and advances, contextual information may change as well, potentially making infrastructure related information become harder to manage in the long term. Presentation composition, taking into account life long experience as well, becomes even harder.

Conclusions and Future Work

It seems that having an integrated, standard, personal "lifelong" user model provides a starting point for personalization in several forms, starting from linking each visit to a particular place with previous visits and, at the extreme, linking each visit with all other aspects modeled about the user. The user model can be stored physically or logically on the visitor's phone, perhaps with parts made available to the museum as needed. An interaction language and protocol should support the request and provision of user personal data, while the semantic differences between the user model ontology and the museum ontology will have to be addressed. The lifelong user model can provide the required infrastructure and information (and the problems of representation and storage of user modeling data are generic). However, it is essential to be able to use user modeling data across-domains and the need to be able to speak the language of every museum (and back) – as far as representing visitors, knowledge, and preferences.

It is worth noting that this paper focused on the individual visitor and on individual lifelong user model. A museum visit is often a social event, where small groups of friends come to explore the museum together. As such, the social context of the visit should be addressed as well, as a natural extension of this initial ideas.

References

Androutsopoulos, J. Oberlander, and V. Karkaletsis. (2007) Source authoring for multilingual generation of personalised object descriptions. Natural Language Engineering. 13 (3): 191–233

Assad, m, D. J. Carmichael, J. Kay, and B. Kummerfeld. (2007) PersonisAD: Distributed, active, scrutable model framework for context-aware services, *Pervasive 07, LNCS,* 55-72. Springer.

Berkovsky, S., Kuflik, T. and Ricci, F. (2008). "Mediation of user models for enhanced personalization in recommender systems". User modeling and User Adapted Interaction Vol. 18 (3), pp 245-286.

Bitgood, S. (2009) Museum Fatigue: A Critical Review. Visitor Studies, 12(2), 93 – 111.

Collins, A, A. Bezerianos, G. McEwan, M. Rittenbruch, R. Wasinger, and J. Kay. Understanding file access mechanisms for embedded ubicomp collaboration interfaces. UbiComp 2009, 135-144, New York, NY, USA, 2009. ACM Press.

Carenini, G., Pianesi, F., Ponzi, M. and Stock. O. (1993) Natural Language Generation and Hypertext Access' Applied Artificial Intelligence, 7(2) 135-164. Hemisphere.

Dim, E. and Kuflik, T. (2010) Social signal processing: detecting small group interaction in leisure activity. In *Procs Internat'l Conf on Intelligent User Interfaces,* 309-312. ACM.

Durbin, G. (Ed.) (1996) Developing museum exhibitions for lifelong learning. Stationery Office, London.

Falk, J. -H. & Dierking, L. D. (1992) *The museum experience*. Washington D.C.: Whalesback.

Falk, J. (2009) Identity and The museum visit experience. Left Coast Press.

Fleck M. Frid M. Kindberg T O'Brien-Strain E. Rajani R. and Spasojevic M. (2002). From Informing to Remembering: Ubiquitous Systems in Interactive Museums, Pervasive computing, 1, 2, 13-21.

Garzotto F. Cinotti T. Pigozzi M. (2003) “Designing multi-channel web frameworks for cultural tourism applications: the MUSE case study”, Selected Papers from Museums and the Web, Toronto Archives & Museums Informatics pp. 239-354

Gerber S., Pink G, Fry M, Kay J, Kummerfeld R, Wasinger R. “PersonisJ: mobile, client-side user modelling”, User Modeling, Adaptation and Personalization (UMAP 2010), to appear.

Hatala, M. and Wakkary, R.: (2005) User Modeling and Semantic Technologies in Support of a Tangible Interface. *Journal of User Modeling and User Adapted Interaction* 15(3-4), Kluwer, 339-380.

Heckmann, D. and Kruger, A. (2003) "A User Modeling Markup Language (UserML) for Ubiquitous Computing". In *9th International Conference on User Modeling (UM'2003)*, Pittsburgh, June, pp. 393-397

Kay, J. (2008) Pervasive personalised lifelong learning and remembering, with user model foundations. IEEE Trans on Learning Technologies, 1(4):215-228.

Kay, J., and B. Kummerfeld. (2009) editors. Proceeding of the Lifelong User Modelling Wkshop at UMAP'09, Trento, Italy http://www.cs.usyd.edu.au/~llum/2009_UMAP09_llum/

Kuflik, T., Sheidin, J., Jbara, S., Goren-Bar, D., Soffer, P., Stock, O. and Zancanaro. M. (2007) Supporting Small Groups in the Museum by Context-Aware Communication Services. IUI Honolulu, Hawaii, USA, pp. 305-308.

Kuflik, T. and Poteriaykina, K. User Model on a Key. In Proceedings of Hypertext 2009, Torino, Italy, June 29th-July 1st 2009 pp. 371-372.

Not, E., Petrelli, D., Sarini, M., Stock, O., Strapparava, C., Zancanaro, M. (1998) Hypernavigation in the physical space: adapting presentations to the user and to the situational context. New Rev. Hypermedia Multimedia J. 4, 33–45

Oberlander, J., O’Donnell, M., Knott, A., Mellish, C. (1998) Conversation in the museum: experiments in dynamic hypermedia with the intelligent labelling explorer. New Rev. Hypermedia Multimedia J. **4**, 11–32

Oppermann, R. and Specht, M. (1999). A Nomadic Information System for Adaptive Exhibition Guidance. Archives and Museum Informatics 13(2): 127–138.

Sarini, M. and Strapparava, C. (1998) Building a User Model for a Museum Exploration and Information-Providing Adaptive System, Procs 2nd Wkshop on Adaptive Hypertext and Hypermedia, Hypertext'98, Pittsburgh, USA, 20-24.

Stock, O,. Zancanaro, M., Busetta, P., Callaway, C., Krüger, A., Kruppa, M.,Kuflik, T., Not E. and Rocchi, C. (2007) Adaptive, Intelligent Presentation of Information for the Museum Visitor in PEACH. *User modeling and User Adapted Interaction.* 17(3), pp 257-304.

Vinciarelli, A., Maja Pantic, Herve Bourlard. (2009) Social signal processing: Survey of an emerging domain, Image and Vision Computing, 27(12) Visual and multimodal analysis of human spontaneous behaviour, 1743-1759

Wang, Y., N. Stash, L. Aroyo, P. Gorgels, L. Rutledge, and G. Schreiber. (2008a). Recommendations based on semantically-enriched museum collections. Journal of Web Semantics: Semantic Web Challenge 6(4): 43.50.

Wang, Y., Aroyo, L., Stash, N., Sambeek, R., Schuurmans Y., Schreiber, G. and Gorgels, P. (2009) Cultivating Personalized Museum Tours Online and On-Site, Journal of Interdisciplinary Science Reviews, Vol. 34, No. 2, Pages 141-156, June, 2009.

Zimmermann, A., Specht, M., Lorenz, A. (2005) Personalization and Context Management. User Modeling and User-Adapted Interaction 15(3-4) 275-302

Modeling Town Visitors Using Features based on the Real World and the Web Information

Junichiro Mori[1], Hitoshi Koshiba[2], Kenro Aihara[2], and Hideaki Takeda[2]

[1] The University of Tokyo, Tokyo, JAPAN
jmori@ipr-ctr.t.u-tokyo
[2] National Institute of Informatics, Tokyo, JAPAN
hkoshiba,aihara,takeda@nii.ac.jp

Abstract. Recent mobile and sensor technologies have enabled acquisition of information about users from their daily living environments. In addition, users' contents and activities on the web provide rich information about their daily life. Given the wide range of user information, it is important to identify relevant features of the user information for automatically constructing a user model of daily life. As described herein, we propose a method that predicts a user model using features based on users' offline behavior, their environments, their online contents and activities. In particular, we specifically examine the domain of modeling a town visitor, where user information of various kinds is available. We capture the visitors' information using various sensors. We also obtain additional information from their blog. Based on the visitors' information from the offline and the online sources, we design features that characterize the visitors in the town. We use those features using a machine learning approach to predict visitor attributes such as age, gender, occupation, and interest. Our results show that our designed features perform with the accuracy of 70–90% for prediction.

1 Introduction

Recent mobile and sensor technologies have enabled acquisition of information about users from their daily living environments. Several studies in the context of ubiquitous and pervasive user modeling have addressed the issue of constructing and representing a user model using such user information from the real world [5]. In addition to the information in the real world, the contents and the activities that the users generate on several web sources such as home pages, blogs, and social media provide rich information about their daily life. Using such user information from the web, several methods for modeling a user have been proposed [2].

There exists a wide range of large amounts of user information available in the real world and the web. Therefore, it is currently possible to collect such information from multiple and heterogeneous sources and construct a user model. In contrast to previous user modeling tasks that have focused on features within a limited and closed domain such as a desktop [7], a museum [1], a classroom [4], and a shop [8], this is a task of user modeling in a rather open and public domain, along the side of "life long" user modeling [6], with a wide range of user information in daily life, which requires a new approach. Thereby, for constructing a user model, it is important to identify relevant features among user information of various kinds from multiple and heterogeneous sources.

As described in this paper, we propose a method that predicts a user model using features based on users' behavior and their environments in the real world and their online contents and activities on the web. In particular, we specifically examine modeling a town visitor as a rather open public domains in which user information of various kinds is available. We conducted an experiment to obtain empirical data from about 200 visitors to Jiyugaoka, a popular town in Tokyo. We capture the visitors' information in the town using various sensors. Additionally, we obtain subjective and social information about the visitors from their blog. Based on the visitors' information from the offline and the online sources, we design features that characterize the visitors in the town. We employ those features using a machine-learning approach to predict visitor attributes such as age, gender, occupation, and interests. Our results show that our designed features perform with accuracy of 70–90% for prediction. Our contributions are two-fold: First, we collect the users' daily life information of various kinds from the real world and the web. Second, we develop a method to predict a user model automatically using such information, which is useful for several applications of ubiquitous user modeling and adaptive hypermedia.

This paper is organized as follows. We introduce our experiment for collecting visitor information from the town and the web in the next section. The proposed method to predict a visitor model using the collected information is explained in Section 3. Analyses of the results are made in Section 4. Finally, we conclude the paper in Section 5.

2 Visitor Data Acquisition

2.1 Data from the real world

To collect town visitor data, we conducted an experiment in Jiyugaoka from January 17 to February 8, 2009. The town is popular in Tokyo as both a residential and shopping area. In all, 196 men and women participants in their 20s–50s participated in the experiment along with 52 shops. The participants were provided an Integrated Circuit card (IC card) and were instructed to walk and visit shops in Jiyugaoka freely according to their interests.

Various shops such as restaurants, bars, grocery stores, hair salons, and fitness club were enrolled in the experiment. We installed a "sensor box" in each shop. The sensor box is about 30 cm long and 20 cm wide. A personal computer is embedded in the sensor box to record the data of shop visitors from the following sensor devices.

IC card reader. An IC card reader is embedded at the front center of the sensor box. It detects a visit of participants at a shop. Whenever visiting a shop, a participant holds her provided IC card over the card reader on the sensor box of the shop so that her visit to the shop is recorded. The record is also used to generate a blog template in our online system, as described below.

Microphone. A microphone is embedded at the upper front left of the sensor box. It detects a sound pressure level around the sensor box. The sound pressure is recorded as a measure of noise in a shop.

Pyroelectric sensor. A pyroelectric sensor is embedded at the front bottom of the sensor box. It detects the amount of traffic of people by counting how often people are walking around the sensor box. The counting is recorded as the measure of congestion in a shop.

Together with these data, we also recorded the data of date, weather, and temperature when a participant visits a shop. All recorded data are sent to the data server via a mobile phone that is embedded in the sensor box. With data collected from the sensor boxes, it is possible to acquire information about town visitors in terms of the shops that they visited and the environments of their visits such as noise, congestion, date, weather, and temperature.

2.2 Data from the Web

To obtain participants' online behaviors in addition to the data from the real world, we provided them with an online system during the experiment. The system enables participants to record and share their experiences in the town and to enhance communication among themselves. To this end, the system generates a blog template that is automatically created based on one's record of visits at shops. After visiting the town during the experiment, the participants can freely edit their blog template and create a blog entry about their visits to shops. The created blog entry becomes shared with other participants so that they can read others' blog entries and comment on them. For each participant, we record his or her blog entries in the form of text information. We also record his or her activities such as accessing, reading, and commenting on the blog.

3 Learning a Visitor Model

In this section, we propose our method to predict a visitor model consisting of several attributes. The visitor model is predicted using a machine learning approach with features based on visitors' data from the real world and the web, as described in a previous section. We first describe our visitor model to be predicted. Then, we explain our feature design to be utilized in a machine learning method.

3.1 Visitor Model

Table 1 shows our visitor model that we define using six attributes (*age*, *gender*, *marital status*, *residential area*, *occupation*, and *interest*). In relation to existing user models, *age*, *gender*, *marital status*, *residential area*, *occupation* correspond to demographics and *interest* corresponds to interests. These attributes were obtained from the questionnaire that each participant filled out before the experiment. For *interest*, we asked "Answer shop categories that you often visit in Jiyugaoka" and allowed multiple answers only for this question.

3.2 Feature design

We design the following features for predicting our visitor model. The basic idea behind the feature design is that there should be particular characterization of participants'

Table 1. List of attributes and values in a visitor model

attribute	value (number of instances / 196)
age	20s (51), 30s (96), 40s (38), 50s (11)
gender	men (46), women (150)
marital status	married (110), single (86)
residential area	Tokyo (100), Kanagawa (90), other (6)
occupation	worker (121), housekeeper (38), student (15, %), other (20)
interest	Western restaurant (93), Chinese restaurant (24), food (44), bar (15), cafe (161), fashion (122), cosmetic (48), commodity (104), culture (39), service (22)

Table 2. List of features

source of data	type of data	feature	number of feature
real world	place	**shop**	52 (string)
		shop category	10 (string)
	environment	**noise**	1×52 (integer)
		congestion	1×52 (integer)
		heat	1×52 (integer)
		weather	3×52 (string)
		temperature	1×52 (integer)
	time	**day**	7×52 (string)
Web	activity	**blog entry**	1 (integer)
		total access	1 (integer)
		total comment	1 (integer)
	communication	**blog access**	200 (integer)
		blog comment	200 (integer)
	content	**blog word**	5,362 (string)
total			6,555

offline and online behaviors and their surrounding environments depending on their attributes. For example, young female students would tend to visit the town on weekdays and spend some time at busy fashion shops and cafe. They would be also active online by creating a lot of contents and frequently communicating with each other. Accordingly, we design our features based on data of six types from the real world and the web, which include place, environment, day, online activity, online communication, online content as presented in Table 2.

shop portrays a visit at a shop by a participant. The visit is based on data from the participant's holding the IC card over the shop's sensor box. **shop** includes the name of 52 shops that were enrolled in our experiment. **shop category** includes 10 categories (Western restaurant, Chinese restaurant, food, bar, cafe, fashion, cosmetic, commodity, culture, and service) that the shops fall into. The categories correspond to the items of question for *interest* of a visitor model.

noise shows the level of noise of a shop, which is based on the sound pressure detected using a microphone. **congestion** shows the level of congestion of a shop. It is based on the amount of traffic of people detected by a pyroelectric sensor. **heat** shows the level of heat of a shop, which is based on the combination of outputs from three sensors, a IC card reader, a microphone, and a pyroelectric sensor, of a sensor box. In fact, It becomes higher when the outputs of the sensors become higher.

weather portrays the prevailing weather when a participant visits a shop. It includes three weather types: sunny, rainy, and cloudy. **temperature** shows the temperature when a participant visits a shop. **day** shows the day of the week when a participant visits a shop.

For **noise**, **congestion**, **heat**, **weather**, **temperature**, and **day**, we obtain these features for each shop. Therefore, the number of respective features expands to the numbers of shop.

blog entry shows the total number of a participant's blog entries. **total access** shows the total number of a participant's accesses to others' blog entries. **total comment** shows the total number of a participant's comments on others' blog entries. **blog access** shows the number of a participant's accesses to respective participants' blog. **blog comment** shows the number of a participant's comments on respective participants' blog. **blog word** portrays a list of words that are included in a participant's blog entries. The words are based on the result of the part-of-speech (POS) tagging applied to the blog entries. We extracted 5,362 nouns, verbs, adjectives, and adverbs as a result of POS tagging. Several weighting functions such as TFIDF and co-occurrence is applicable to the words. For simplicity, we only use frequency-based weighting after normalizing the words based on the number of blog entries.

For **noise**, **congestion**, **heat**, **weather**, **temperature**, and **day**, we prepared these features for each shop. Therefore, one learning instance, one participant in our case, is represented with features of 14 kinds. These 14 types of feature consist of a 6,555 dimensional feature vector as shown Table 2. For each participant, we created a feature vector based on visits to the town and that individual user's online activities including editing, accessing, and commenting on blog entries. Some participants visited the town several times during the experiment. Therefore, features related to their visits are normalized with times of their visits. Regarding features related to blogs, **blog access** and **blog comment** are normalized with **total access** and **total comment** respectively.

3.3 Prediction

We obtained a feature set with data of 196 participants. Given the visitor model and the feature set, our task is now to predict each attribute in the visitor model (described in Sect. 3.1) using the feature set (described in Sect. 3.2). For each attribute in the visitor model, we train a learner that predicts the attribute (*gender*, *age*, *interest*, etc.) given a set of feature values (**shop**, **weather**, **blog word**, etc.) corresponding to the certain value of the attribute (*gender*: men or women, *age*: 20s, 30s, 40s, or 50s, etc.). For example, given a user's activities and its environment in the town together with his or her online activities on the blog, the trained learner would predict his or her gender, age, or interest in the visitor model

As a learner, we use a support vector machine (SVM) [9]. We use a radius basis function (RBF) kernel, which performs well in our preliminary experiments. For each

Table 3. Performances of prediction for respective attributes

attribute	feature type 1	feature type 2	feature type 3
gender	73.912	85.869	88.043
marital status	69.767	85.465	86.046
age	56.375	69.127	70.469
resindece area	66.666	82.812	85.416
occupation	74.0	84.0	86.0
interest			
1. Western restaurant	69.892	79.569	81.182
2. Chinese restaurant	81.25	91.666	93.75
3. food	78.409	86.363	87.5
4. bar	86.666	90.0	93.333
5. cafe	72.857	87.142	87.142
6. fashion	75.0	81.756	81.756
7. cosmetic	77.083	85.416	88.541
8. commodity	69.565	83.152	81.521
9. culture	74.358	83.333	84.615
10. service	70.454	88.636	90.909
mean	73.083	84.287	85.748
stddev	6.882	5.254	5.706

feature type 1: real-world-oriented features
shop, **shop category**, **noise**, **congestion**, **heat**, **weather**, **temperature**, **day**
feature type 2: web-oriented features
blog entry, **total access**, **total comment**, **blog word**, **blog access**, **blog comment**
feature type 3: type 1 and 2 (complete features)

attribute, we solve the two-class or multi-class problem depending the number of attribute values.

To train the learner and evaluate its performance, we splitted our data into five data set. We use four data set for training the learner and test the learner using the remaining one data set. With this data set, we evaluate the performance using five-fold cross validation. We compare the performances of three feature types. The first feature type (type 1) includes features that are based on user information from the real world including **shop**, **shop category**, **noise**, **congestion**. The second feature type (type 2) includes features that are based on user information from the web including **blog entry**, **total access**, **total comment**, **blog word**, **blog access**, and **blog comment**. The third one includes all the features.

4 Evaluation and Results

Performances of the learner that is trained with respective types of feature are presented in Table 3 with prediction accuracy for respective attributes. The performance of the learner varies depending on the attribute to be predicted. Depending on the attribute,

Table 4. Distribution of selected features for respective attributes

feature	attribute							
	1	2	3	4	5	6	mean	stdev
Blog word	71.86	68.62	73.77	82.37	75.88	69.32	73.63	5.06
day	6.06	11.37	6.01	6.08	8.58	7.90	7.66	2.11
Blog access	0.43	3.1	3.64	5.44	4.37	4.06	3.50	1.69
weather	3.46	3.44	3.09	1.28	2.62	3.26	2.85	0.83
temperature	4.32	3.44	2.55	0.96	1.57	2.69	2.58	1.21
noise	3.03	2.41	2.36	1.60	1.40	2.95	2.29	0.67
shop	3.89	2.06	2.36	0.64	1.75	2.89	2.26	1.09
heat	3.89	2.06	2.67	0.64	1.22	2.96	2.24	1.18
congestion	2.16	1.72	2.00	0.32	0.87	2.94	1.66	0.94
shop category	0.86	0.68	0.91	0.32	0.52	0.65	0.65	0.21
Blog comment	0	0	0.36	0	0.17	0.02	0.09	0.14
total comment	0	0.34	0.18	0	0	0.07	0.09	0.13
Blog entry	0	0	0.18	0	0	0.07	0.04	0.07
total access	0	0	0	0	0	0.02	0	0

attribute: 1.*gender*, 2.*marital status*, 3.*age*, 4.*resindetial area*, 5.*occupation*, 6.*interest* (average of ten items)

the performance varies as much as 20 points, which indicates that some attributes are predicted and that others are difficult to predict. On average, the accuracy with the feature type 1 is about 73.08%, about 84.28% with the feature type 2, and about 85.74% with the feature type 3. Because the learner with feature type 3 performs well for every attribute, the t-test (two-tailed, assuming equal variances) demonstrates its significant performance as opposed to feature type 1 (p = 4.869e-9) and type 2 (p = 7.281e-4). However, the performance with the feature type 2 is comparable with the feature type 3 for several attributes, which suggests that the features based on the user information from the web are important for our prediction.

To investigate which particular features are effective for predicting respective attributes, we further analyze the results of feature selection. To determine a proper set of features among our many features, we use a feature selection strategy [3]. Table 4 presents the distribution of selected features for respective attributes. Several features based on information from users' environments such as **day**, **weather**, and **temperature** are selected for the prediction. Other environment-oriented features such as **congestion** and **heat** are less influential, which suggests that the features (**day**, **weather**, and **temperature**) based on rather static and global environment affect the prediction than the features (**congestion**, **heat**, and **noise**) based on the static and local environment. Regarding the features based on users' behavior information such as **shop** and **shop category**, they are less influential for prediction than the environment-oriented features, which suggests that the environment surrounding a user is an important factor for modeling a user in an open and public domain.

Among web-oriented features, **blog access** is relatively influential for the prediction, which shows that users' patterns of online information access differs depending on their attributes. Apparently, **blog word** has a strong effect on the prediction among all the features. In fact, if we look at the words of selected **blog word** features carefully, then we can find particular words for respective attributes. For example, “Valentine’s Day” is selected as one distinguishing words for *gender*. The words “baby stroller”, “husband”, and “child-friendly” are selected to predict users’ *marital status*. For *job*, words such as “company”, “coming home”, and “Sunday” are selected. Several words that express users’ particular preferences are included for respective items of *interest*. For *interest* in cafe, words such as “bread”, “coffee”, and “wine” are selected; for *interest* in fashion, the words such as “boots”, “clerk”, and “beauty” are selected. As these examples show, **blog word** is a useful and intuitive feature that is might be helpful for prediction because it represents particular subjective sentiments and preferences depending on users’ attributes.

5 Conclusion and Future work

This paper has presented a proposal of a method to predict a user model using features based on user information of various kinds from the real world and the web. In particular, we described our specific examination of modeling a town visitor. Our results demonstrate that our designed features perform with the accuracy of 70–90% for prediction. The results also reveal that the features based on users’ online contents and activities with the complements of environment-oriented features are important for predicting a user model. In future work, we will use more complicated features that characterize various types of offline and online behavior patterns for modeling a user.

References

1. Bohnert, F., Zukerman, I.: Non-intrusive personalisation of the museum experience. In: Proc. 1st and 17th International Conference on User Modelling, Adaptation, and Personalization (UMAP 2009). pp. 197–209 (2009)
2. Brusilovsky, P., Kobsa, A., Nejdl, W. (eds.): The Adaptive Web: Methods and Strategies of Web Personalization. Springer Berlin / Heidelberg (2007)
3. Chen, Y.W., Lin, C.J.: Combining SVMs with various feature selection strategies. Springer (2006)
4. Cooper, D.G., Arroyo, I., Woolf, B.P., Muldner, K., Burleson, W., Christopherson, R.: Predicting customer models using behavior-based features in shops. In: Proc. 1st and 17th International Conference on User Modelling, Adaptation, and Personalization (UMAP 2009). pp. 30–41 (2009)
5. Heckmann, D.: Ubiquitous Use Modeling. Ph.d thesis, University of Saarland (2005)
6. Kay, J.: Lifelong leaner modeling for lifelong personalized pervasive learning. IEEE Transactions on Learning Technologies 1(4), 215–228 (2009)
7. Kleek, M.V., Shrobe, H.E.: A practical activity capture framework for personal, lifetime user modeling. In: User Modeling. pp. 298–302 (2007)
8. Mori, J., Matsuo, Y., Koshiba, H., Aihara, K., Takeda, H.: Predicting customer models using behavior-based features in shops. In: Proc. 1st and 17th International Conference on User Modelling, Adaptation, and Personalization (UMAP 2009). pp. 126–137 (2009)
9. Vapnik, V.: The Nature of Statistical Learning Theory. Springer-Verlag (1995)

Modeling Health Problems of Elderly to Support their Independent Living

Bogdan Pogorelc[1,2], Matjaž Gams[1,2]

[1] Jožef Stefan Institute, Department of Intelligent Systems, Jamova 39, Ljubljana, Slovenia
[2] Špica International d. o. o., Pot k sejmišču 33, Ljubljana, Slovenia
{bogdan.pogorelc, matjaz.gams}@ijs.si

Abstract. We propose a system for modeling and recognition of gait patterns related to health problems of elderly to support their independent living. Gait of elderly is captured with motion capture system, which consists of tags attached to the body and sensors situated in the apartment. Position of the tags is acquired by the sensors and the resulting time series of position coordinates are analyzed with machine learning algorithms in order to recognize the specific health problem. We used medically inspired features for training a machine learning classifier that classifies the user's gait into: i) normal, ii) with hemiplegia, iii) with Parkinson's disease, iv) with pain in the back and v) with pain in the leg. Experimental results showed that decision tree classifier was able to reach only 95% of classification accuracy, and random forest of over 99%, using only 8 tags with 0-20 mm noise.

Keywords: Health problems detection, human motion analysis, gait analysis, machine learning, data mining, human locomotion, ambient intelligence.

1 Introduction

The rapid aging of population is becoming increasingly difficult issue in the developed countries, since the society's capacity is becoming overwhelmed with taking care of its elderly members [1].

Elderly want to live at their homes for as long as possible. However, in older age they become more prone to specific health problems. When living alone, nobody could detect changes in behavior potentially being health problem and call for medical help. To prevent such dangerous situations, we propose system for automatic ubiquitous elderly health care, for which we developed techniques for recognition of common health problems manifesting in gait of elderly. In case the system would recognize the health problem, it would automatically notify the physician and show him/her the explanation of the automatic diagnosis in form of visualization of the kinematic model. Therefore, elderly would get constant health care at their homes and physicians would be less occupied with work, however they would still have the possibility to confirm/reject automatic diagnose. In this case elderly would gain constant (ubiquitous) monitoring, providing them more safety and confidence in living at their homes.

The common way of diagnosing health problems is by physician, when the elderly goes to the medical examination, which happens too rarely for many elderly. Early detection of health problems would provide much better chances for recovery. The proposed system is not anticipated to replace the physician, but it would provide early warning and suspected health problem to physician, that he/she would then make detailed examination of the elderly and confirm/reject automatic diagnosis.

The need for such system was expressed by European Commision through the 7th Framework Programme (FP7) and Ambient Assisted Living (AAL) as preferential area. Solutions, described in the paper were also partially implemented in FP7 project Confidence [2, 3], which intends to provide confidence for the elderly who want to live at their homes.

The target health problems for automatic recognition are: hemiplegia (usually the result of stroke), Parkinson's disease, pain in the leg and pain in the back. The gait of the user is captured with the motion capture system, which consists of the tags attached to the body and the sensors situated in the apartment. The position of the tags is acquired by the sensors and the resulting time series of position coordinates are analyzed with machine learning algorithms in order to recognize the specific health problem.

Concerning privacy issues, investigation of the FP7 project Confidence [3] showed that elderly do not want to be monitored by cameras which record video, which is solved in our system. Moreover, we investigated the classification accuracy achievable using various numbers/placements of tags on the user's body and various amounts of noise in tag coordinates. Tag placement must achieve a trade-off between usability and technical requirements – the users prefer as few tags as possible, but too few tags cannot ensure sufficient accuracy. Both the finding regarding noise and tag placement can affect motion-capture system selection and further development and applications of care systems for elderly.

For the automatic recognition of the movement (gait) pattern, the movement must first be captured. For this purpose many types of motion capture devices exist. Widely used are inertial motion capture systems composed of accelerometers or gyro sensors [4-6]. The second widely used approach uses machine vision for the reconstruction of the human body movement [20, 22]. The third approach uses cameras in combination with tags attached to the body. Usually infra-red (IR) cameras are used and the body posture is reconstructed from the position of the retroreflective tags [14], as in our approach. There also exist some specific measurement devices for the recognition of tremor – a symptom in Parkinson's disease, but not in hemiplegia, pain in the leg and pain in the back. Tremor can be acquired with variety of measuring approaches, including sensors for a measurement of the angle of the joint deflection in tremor-type joint movements [7] and with electromyography [8].

We did not address the recognition of the activities of daily living, such as walking, sitting, lying, etc., and the detection of falling, which has already been solved [3, 17]. We were focused on solving a more challenging task, which is the recognition of gait-related health problems.

In works related to health problems recognition [11, 12, 13], physicians usually diagnose health problems which manifest in gait just by manually observing the user's gait. However, this approach cannot provide constant real-time observation of the elderly at home, for fast recognition of changes in movement (gait), indicating some health problem. This is also the case in [23], where a system for long-term monitoring of the gait in Parkinson's disease is presented. The characteristics of every stride taken were acquired using a lightweight ankle-mounted sensor array that saved the data into a small pocket PC. The work [7] presents sensors for the measurement of the angle of joint deflection in tremor-type joint movements, which can also be used to assess Parkinson's disease. Just like the system described in [23], it has major drawback in comparison to our approach, because it cannot automatically recognize Parkinson's disease or any other health problem.

Using similar motion-capture system as that in our approach, the automatic distinguishing between health problems such as hemiplegia and diplegia is presented in [16]. Hemiplegia and diplegia are common states after cerebral palsy and sometimes also after stroke. Hemiplegia is (partial) paralysis of a person on one side of the body and diplegia is (partial) paralysis on both sides. A classification accuracy of 92.5% was reported. This was achieved with Self-Organizing Maps, whose features were wavelet-transformed gait characteristics such as walking speed and stride length.

An important part of the research presented in this paper is the study of the impact of the placement of tags on the user's body and the amount of noise in tag coordinates on the classification accuracy. The closest work in this respect that we are aware of investigated the placement of accelerometers for fall detection [10, 15]. Their finding was that the head provides optimal accuracy, but is impractical, the wrist is not appropriate, and the waist is a good option.

2 Methods and materials

The specific health problems for the development of our health problem recognition system were suggested by the medical expert based on the incidence in the elderly aged 65+, the medical significance and the feasibility of their recognition from the observed subjects' movements. The following four health problems were chosen as the most appropriate: *Parkinson's disease*, *hemiplegia*, *pain in the leg* and *pain in the back.* The fifth chosen health state was *normal* (healthy) and was used as a reference.

A physician usually diagnoses target health problems while observing a patient's gait (i.e. posture and the walking pattern). Since the gaits of patients with the observed five health states look similar to each other, a physician needs to pay attention to many details to recognize the health state [11, 12, 13]. For the task of the automatic health-state recognition we proposed and tested features that are based on the tag locations, for 12 tags, placed on the shoulders, elbows, wrists, hips, knees and ankles of the elderly. Some of the proposed features, which are used for modeling using the machine learning methods, are listed as follows:

- Absolute difference between i) average distance between right elbow and right hip and ii) average distance between right wrist and left hip.
- Average angle of the elbow.
- Quotient between maximal angle of the left knee and maximal angle of the right knee.

- Quotient between i) difference between maximal and minimal height of left ankle and ii) maximal and minimal height of right ankle.
- Absolute difference between i) difference between maximal and minimal speed (magnitude of velocity) of the left ankle and ii) difference between maximal and minimal speed of the right ankle.
- Average speed (magnitude of velocity) of the wrist.
- Frequency of angle of the elbow passing average angle of the elbow.
- Average angle between i) vector between right shoulder and right hip and ii) vector between right shoulder and right wrist.
- Difference between average height of the right shoulder and average height of the left shoulder.

For developing predictive model for automatic recognition of health problems in the subjects yet to be observed, we employed supervised learning methods from the field of machine learning. In supervised learning, a training data set of already labeled subjects (i.e. classified into one of the target five classes) is used to construct a model, which is later used to predict the class of the subjects for which we wish to detect the health problem. Our task was therefore to classify the recordings of walking into five classes: four with selected health problems (classes *hemiplegia*, *parkinson*, *pain-leg*, *pain-back*) and one without it (*normal*).

Data for the evaluation of the proposed approach was collected by recording the gaits of test subjects with particular walking patterns. The final data set of 141 recordings consisted of: i) 25 recordings of normal walking, ii) 45 recordings of walking with hemiplegia, iii) 25 recordings of walking with Parkinson's disease, iv) 25 recordings of walking with a limp due to a pain in the leg, v) 21 recordings of walking with a limp due to a pain in the back.

The recordings consisted of the position coordinates for the 12 tags worn on the body, sampled with 60 Hz. The tag coordinates were acquired with Smart IR motion capture system with 0.5 mm standard deviation of noise. For each subject, the locations of the sensor tags were recorded in a session which lasted 5-8 seconds from which a vector of 13 proposed features was computed. These learning examples were labeled with the type of the representing health problem, yielding the final data on which the classifier was trained.

3 Experiments and Results

In our experimental work we focused on analyzing the classification accuracies of model, built using the machine learning methods. The experimental classification accuracies were obtained using stratified 10-fold cross validation. We used the decision tree and random forest machine learning algorithms implemented in Weka [21].

The 10-fold cross-validation of decision tree and random forest resulted in classification accuracy of 90.1% and 99.3%, respectively.

Table 1. Confusion matrices of decision tree (left) and random forest classifier (right), where H=hemiplegia, L=pain in the leg, N=normal (healthy) walking, P=Parkinson's disease and B=Pain in the back. Numbers denote numbers of the classified examples.

		classified as				
		H	**L**	**N**	**P**	**B**
true class	**H**	40	0	1	4	0
	L	2	23	0	0	0
	N	2	0	23	0	0
	P	5	0	0	20	0
	B	0	0	0	0	21

		classified as				
		H	**L**	**N**	**P**	**B**
true class	**H**	45	0	0	0	0
	L	0	25	0	0	0
	N	0	0	25	0	0
	P	1	0	0	24	0
	B	0	0	0	0	21

Table 1 shows the confusion matrices, i.e. how many examples of a certain true class (in rows) are classified in one of possible five classes (in columns). We can use confusion matrices for three purposes:

- We can see how many false positives (false alarms) can be expected using these classifiers. When in practice the system would report false alarm, e.g., normal walking is classified as some health problem, ambulance could erroneously drive to pick up the elderly which would cause unnecessary costs. The only false positives occurred in the case of decision tree, where 2 of 25 examples of normal walking were classified as hemiplegia.
- We can see how many false negatives can be expected using these classifiers. False negatives could mean potentially risky situation for the elderly, as his/her health problem would not be recognized automatically. The only false negatives occurred in the case of decision tree, where 1 of 45 examples of hemiplegia was classified as normal walking.
- We can identify between which health states (classes) the most errors (misclassifications) occurs. Consequently, we can add additional features to help distinguish between those particular classes. The most misclassifications were observed between hemiplegia and Parkinson's disease; i.e., 5 of 25 examples and 1 of 25 examples of Parkinson's disease were wrongly classified as hemiplegia for decision tree and random forest, respectively.

The results of random forest were overall better than the results of decision tree. They show that in the proposed approach false positives/negatives are very rare, i.e., the system would provide great safety and confidence for elderly and at the same time it would not cause much unnecessary ambulance costs.

3.1 Dependence of the Classification Accuracy on the Tag Placement and Noise Level

To test the robustness of the approach, we added Gaussian noise with varying standard deviation (and zero mean) to the raw coordinates. The standard deviation of noise was varied from 0 mm to 50 mm in steps of 5 mm. As a preprocessing step, a Kalman filter was used to smooth the potentially unrealistic difference between the positions of two consecutive time samples, caused by the addition of Gaussian noise to the captured positions [18].

Since wearing the full complement of 12 tags may be annoying to the user, we investigated ways to reduce the number of tags. We started with all 12 tags and removed them in the order that retained the largest number of the features when decreasing the number of tags by one. Consequently, the "best" tag placement for each number of the 12 tags was obtained.

Fig. 1 shows the dependence of the classification accuracy (CA) on the tag placement and the noise level. We can observe a variation of the noise in the standard deviation from 0 to 50 mm on horizontal axis and the best tag placement for each number of tags from 12 to 1 tag on the vertical axis. Each curve of different color and shape (e.g. doted, dashed) connects points of the particular classification accuracy.

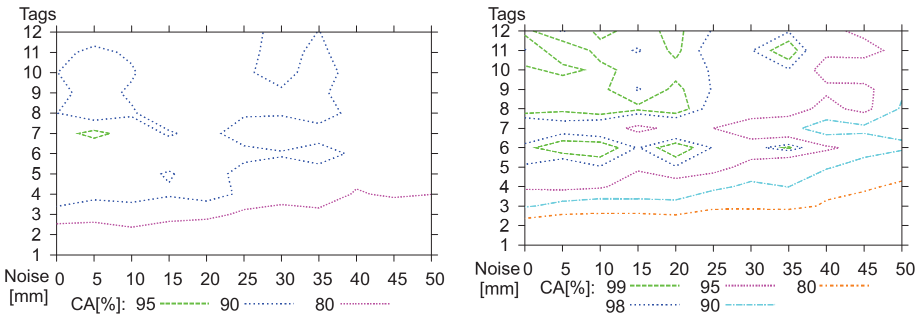

Fig. 1. Classification accuracy with respect to the number of tags and the noise level for the recognition of the target health problems using the decision tree (left) and random forest classifier (right).

Fig. 1 (left) illustrates the classification accuracies for the decision-tree classifier, the highest classification accuracy over 95% is reached only for 7 tags and 5 mm noise. To surpass the classification accuracy of 90% we need at least 4-6 tags, depending on the noise. To exceed an accuracy of 80% at least 3 tags are required for 0-20 mm and 4 tags for 25-50 mm noise.

Fig. 1 (right) shows the classification accuracies for the random forest classifier. It needed only 8 tags with noise 0-20 mm to reach the classification accuracy of over 99% except for the "islands" of insignificantly lower accuracies. To reach the accuracy of over 95%, 4-6 tags were required, depending on the noise. To exceed the boundary of 90%, 3-6 tags were needed. To exceed the accuracy of 80%, at least 3 tags for 0-35 mm, 4 tags for 35-45 mm and 5 tags for 50 mm were needed.

To conclude, random forest overall achieved higher classification accuracies than decision tree classifier. Decision tree exceeded only 95%, while random forest exceeded 99% classification accuracy. Although decision tree achieved lower accuracies than random forest, it has important advantage. Results of its classification are easily interpretable; i.e., if in practice the system would recognize some health state, the physician could check, why the system thinks it is certain health state.

4 Conclusion

A system for modeling and recognition of gait patterns related to health problems of elderly to support their independent living is proposed in this paper.

The results show that in the initial setting (no noise, all tags) decision tree and random forest algorithm achieved an average classification accuracies of 90.1% and 99.3%, respectively. False positives/negatives occurred very rarely, i.e., the system would provide great safety and confidence for elderly and at the same time it would not cause much unnecessary ambulance costs.

The results of investigation of the dependence of the classification accuracy on the tag placement and noise level show that the decision tree was able to reach only 95% of classification accuracy. This was the case only with 7 tags and 5 mm noise. On the other hand, random forest was much more accurate since it reached classification accuracy of over 99% with only 8 tags with 0-20 mm noise. However, decision tree has important advantage over random forest, since the results of its classification are easily interpretable; i.e., if in practice the system would recognize some health state, the physician could check, why the system thinks it is certain health state.

In future work, we see potential to upgrade personalization. Elderly would have mobile device with user interface which would allow them to reject warning if their usual activity (such as yoga) is diagnosed as health problem.

There is also possibility of personalization for the elderly, who already have some health problem detected. In this case, their usual movement is recorded to be used as training examples for the machine learning algorithms. If their gait patterns deviate from their usual gait patterns, the system triggers warning for the remote medical center.

Another possibility is automatically controlled rehabilitation, where the user is awarded if his/her movement improves in comparison to the previous state.

Acknowledgments. This work is partially financed by the European Union, the European Social Fund. The authors thank Martin Tomšič, Bojan Nemec and Leon Žlajpah for their help with data acquisition, Anton Gradišek for his medical expertise and Zoran Bosnić and Mitja Luštrek for helpful discussions.

References

1. S. Toyne, Ageing: Europe's growing problem. BBC News. http://news.bbc.co.uk/2/hi/business/2248531.stm, 2009-01-19.
2. M. Luštrek, B. Kaluža, E. Dovgan, B. Pogorelc, and M. Gams, 2009. “Behavior Analysis Based on Coordinates of Body Tags”, Lecture Notes In Computer Science, vol. 5859, 14-23. DOI= http://dx.doi.org/10.1007/978-3-642-05408-2_2
3. Confidence: Ubiquitous Care System to Support Independent Living. http://www.confidence-eu.org.
4. D. Strle, V. Kempe, “MEMS-based inertial systems”, Informacije MIDEM 37(2007)4, pp. 199-209.
5. D. Jurman, M. Jankovec, R. Kamnik, M. Topič, “Inertial and magnetic sensors: The calibration aspect”, Informacije MIDEM 37(2007)2, pp. 67-72.

6. F. Dimic, B. Mušič, R. Osredkar, “An example of an integrated GPS and DR positioning system designed for archeological prospecting”, Informacije MIDEM 38(2008)2, pp. 144-148.
7. S. Ribarič, J. Rozman, “Sensors for measurement of tremor type joint movements”, Informacije MIDEM 37(2007)2, pp. 98-104.
8. J. Trontelj, J. Trontelj and L. Trontelj, “Safety Margin at mammalian neuromuscular junction – an example of the significance of fine time measurements in neurobiology”, Informacije MIDEM 38(2008)3, pp. 155-160.
9. Bourke, A.K., and Lyons, G.M. A threshold-based fall-detection algorithm using a bi-axial gyroscope sensor. Medical Engineering & Physics 30, 1 (2006), 84–90.
10. Bourke, A.K., Scanaill, C.N., Culhane, K.M., O'Brien, J.V., and Lyons, G.M. An optimum accelerometer configuration and simple algorithm for accurately detecting falls. In Proc. BioMed 2006 (2006), 156–160.
11. Craik R., and Oatis C. Gait Analysis: Theory and Application. Mosby-Year Book (1995).
12. Harrison RD. Harrison's principles of internal medicine. New York: McGraw-Hill, 14th edition, 1998.
13. Perry J. Gait Analysis: Normal and Pathological Function. McGraw-Hill, Inc., 1992.
14. eMotion. Smart motion capture system. http://www.emotion3d.com/smart/smart.html.
15. Kangas, M., Konttila, A., Lindgren, P., Winblad, P., and Jamsa, T. Comparison of low-complexity fall detection algorithms for body attached accelerometers. Gait & Posture 28, 2 (2008), 285–291.
16. Lakany, H. Extracting a diagnostic gait signature. Pattern recognition 41 (2008), 1627–1637.
17. Luštrek M., Kaluža B., Fall Detection and Activity Recognition with Machine Learning. Informatica (Slovenia) 33(2): 197-204 (2009).
18. Maybeck, P.S. Stochastic models, estimation, and control. Mathematics in Science and Engineering 141 (1979).
19. Tapia, E.M., Intille, S.S., Haskell, W., Larson, K., Wright, J., King, A., and Friedman, R. Real-time recognition of physical activities and their intensities using wireless accelerometers and a heart rate monitor. In Proc. ISWC 2007 (2007), 37–40.
20. Vishwakarma, V., Mandal, C., and Sura, S. Automatic detection of human fall in video. Lecture Notes in Computer Science 4815 (2007), 616–623.
21. Witten, I.H., and Frank, E. Data Mining: Practical Machine Learning Tools and Techniques (2nd edition). Morgan Kaufmann (2005).
22. Zouba, N., Boulay, B., Bremond, F., and Thonnat, M. Monitoring activities of daily living (ADLs) of elderly based on 3D key human postures. In Proc. ICVW 2008 (2008), 37–50.
23. Moore ST, et al., Long-term monitoring of gait in Parkinson’s disease, Gait Posture (2006), doi:10.1016/j.gaitpost.2006.09.011

Personalising the Museum Experience

Fabian Bohnert

Faculty of Information Technology, Monash University
Clayton, VIC 3800, Australia
fabian.bohnert@monash.edu

Abstract. Visitors to physical museums are often overwhelmed by the vast amount of information available in the space they are exploring, making it difficult to select personally interesting content. In contrast to visits to online museum collections, the selection process is complicated by the facts that (1) it takes time for people to move between exhibits; and (2) exhibitions may be arranged in a way that does not reflect visitors' personal interests, meaning that the interesting exhibits may be scattered throughout the museum. Recent advances in mobile technology and user modelling have enabled computer-based systems that can assist visitors in selecting interesting content. Additionally, such systems can provide visitors with personalised information about exhibits while they are exploring the museum. This paper categorises state-of-the-art technology for personalising visitors' experiences in museums, and discusses current challenges for enabling personalised visitor-support systems.

1 Introduction and Overview

Encouraged by a shift towards more visitor orientation in the 1960s and 1970s, museums have long since evolved from institutions which paid little attention to visitors' needs to places where people go to learn in an enjoyable environment, while seeking information and education about art, science and cultural heritage [1]. Hand in hand with the trend towards visitor focus and engagement goes the provision of differentiated access to information and services tailored to a visitor's specific profile,[1] both in the physical museum and online. This is because viewing personally interesting content encourages visitors' engagement and learning [2, 3]. For example, museums have seen automated self-guided visitor information systems appear and evolve from audio to audio-visual [3], and recently the first attempts have been made to personalise content presentation based on real-time positioning and a visitor's interests, e. g., [4–6].

The research fields of museum informatics, pervasive computing and user modelling have all provided key approaches and technologies to enable personalised museum experiences. They can be categorised as follows.

Types of personalised services. With the goal of providing a more enjoyable experience, personalised visitor-support systems can provide a variety of

[1] The interests and goals of museum visitors depend on a large number of factors, such as visitors' cultural backgrounds, previous experiences and social contexts [2].

services. These services include (1) delivering personalised (multimedia) content about exhibits in the physical museum [5–7], (2) delivering recommendations on personally interesting museum exhibits [8–10], (3) stimulating interaction with the museum environment by linking multimedia content with content in the museum [5], and (4) encouraging social interaction with other museum visitors [11].

Types of user modelling time frame. Personalised visitor-support systems can (1) employ single-visit user models [5–8], (2) access user modelling data that has been acquired before a visit [9], or (3) exploit cross-visit or long-term user models [12]. The challenge for single-visit user models is to adapt to a visitor's interests early during a visit. This adaptation process can be sped up by initialising user models at the beginning of the visit, e. g., by manually bootstrapping the user models [6], or utilising external sources like those from cases (2) and (3).

Types of user model construction. Personalised visitor-support systems for physical museums can employ adaptable user models that require people to explicitly state their interests in some form, e. g., [9]. Alternatively, preferences and interests can be estimated from non-intrusive observations, utilising adaptive user models that do not require explicit visitor input [5–8, 13]. Adaptive visitor-support systems have often primarily updated their user models from visitors' interactions with the system, e. g., [5–7]. Alternatively, user model updates can be based upon non-intrusive observations of visitors' movements through the museum [6, 8, 13, 14].

Types of domain knowledge representation. Visitor-support systems for the museum domain have often used an explicit, a-priori engineered representation of the domain knowledge to enable personalised museum experiences, e. g., [5–7, 9]. Alternative approaches include statistical user modelling techniques, which do not require an explicit domain knowledge representation [8, 13–15]. Such techniques may be advantageous in the context of large museums, where a comprehensive domain knowledge representation may be hard to achieve.

Types of technology. Hardware technologies for personalised visitor-support systems include personal handheld devices such as smartphones [5, 10], embedded computing devices [6], and combinations of the two technologies. For non-intrusive systems based on visitors' movements, these technologies are typically combined with instrumentation for sensing visitors' behaviour (e. g., location and direction) in the museum.

2 Current Challenges

The museum domain (in particular, its physicality) provides specific research challenges for personalised visitor-support systems. Current challenges include the following.

Achieving non-intrusiveness and adaptiveness. While personalised visitor-support systems can initialise their user models by using external user modelling data [12], these user models should also be dynamically updated with the progression of a visit to incorporate additional, site-specific information about visitors. Automatic updates can be achieved by processing non-intrusively obtained sensor information about visitors' movements and behaviour in a museum. User models based on this type of data have already been proposed, e. g., [5, 6, 8, 10, 13, 14], but achieving adaptiveness and context-awareness from location information still poses practical challenges due to the difficulty associated with accurately positioning museum visitors inside museum buildings. Deployment of such systems will require (1) further research in the area of indoor positioning technology to automatically track visitors during their visits, and (2) techniques for linking sensor information with user model input that appropriately consider the impact of measurement noise on the performance of user models (initial research in this area includes [16] and [17]).

Generating personalised content. Many research projects have investigated personalised content delivery systems for museums, e. g., [5–7]. Remaining challenges include achieving increased coherence within personalised presentations, and linking in-situ presentations with direct access to online museum collections.

Generating exhibit recommendations. In contrast to traditional domains for recommender systems, predictions differ from recommendations in physical museums. For example, it may be advisable not to recommend interesting exhibits that are expected to be visited immediately anyway (to avoid annoyance with the system). Similarly, recommendations about interesting exhibits that are far away should be delivered only under extreme circumstances (e. g., if the museum is about to close), in order not to interrupt a visitor's experience. Current solutions for recommending personalised museum tours or exhibit themes do often not consistently consider such factors.

Linking pre-visit, in-situ and post-visit museum experiences. While some researchers have investigated approaches for linking pre-visit, in-situ and post-visit museum experiences, e. g., [5, 9, 18, 19], recent technological developments have generated additional challenges. These challenges include linking in-situ museum experiences with online access to museum collections, and in-situ visitor interaction with online social networking technologies.

Limited computational resources. Handheld devices have limited processing capabilities. Client/server architectures have often been used to carry out computationally expensive operations on a high-capacity server at the backend, but this solution requires wireless connectivity. Server-independent solutions will require resource-efficient software for handheld devices.

3 Final Remarks

This paper provided an assessment of pervasive user modelling and personalisation techniques for personalising the experiences of museum visitors. While previous research projects have already tackled numerous challenges in this domain, many challenges remain and new opportunities have emerged for future research. This paper discussed some of these challenges and opportunities.

References

1. Filippini-Fantoni, S.: Museums with a personal touch. In: Proc. of the 10th EVA London Conf. (EVA-London-03). (2003) 25.1–25.10
2. Falk, J.H.: Identity and the Museum Visitor Experience. Left Coast Press (2009)
3. Tallon, L., Walker, K., eds.: Digital Technologies and the Museum Experience – Handheld Guides and Other Media. AltaMira Press (2008)
4. Proctor, N., Tellis, C.: The state of the art in museum handhelds in 2003. In: Proc. of the 12th Intl. Conf. on Museums and the Web (MW-03). (2003)
5. Stock, O., Zancanaro, M., Busetta, P., Callaway, C., Krüger, A., Kruppa, M., Kuflik, T., Not, E., Rocchi, C.: Adaptive, intelligent presentation of information for the museum visitor in PEACH. UMUAI **18**(3) (2007) 257–304
6. Hatala, M., Wakkary, R.: Ontology-based user modeling in an augmented audio reality system for museums. UMUAI **15**(3-4) (2005) 339–380
7. Petrelli, D., Not, E.: User-centred design of flexible hypermedia for a mobile guide: Reflections on the HyperAudio experience. UMUAI **15**(3-4) (2005) 303–338
8. Bohnert, F., Zukerman, I.: Non-intrusive personalisation of the museum experience. In: Proc. of the 17th Intl. Conf. on User Modeling, Adaptation, and Personalization (UMAP-09). (2009) 197–209
9. Wang, Y., Aroyo, L., Stash, N., Sambeek, R., Schuurmans, Y., Schreiber, G., Gorgels, P.: Cultivating personalized museum tours online and on-site. Interdisciplinary Science Reviews **34**(2) (2009) 141–156
10. van Hage, W.R., Stash, N., Wang, Y., Aroyo, L.: Adaptation step-by-step: Challenges for real-time spatial personalization. In: Proc. of the 2010 Workshop on Pervasive User Modeling and Personalization (PUMP-10 at UMAP-10). (2010)
11. Wakkary, R., Muise, K., Tanenbaum, K., Hatala, M., Kornfeld, L.: Situating approaches to museum guides for families and groups. In: Proc. of the 10th Intl. Cultural Heritage Informatics Meeting (ICHIM-07). (2007)
12. Kuflik, T., Kay, J., Kummerfeld, B.: Life-long personalised museum experiences. In: Proc. of the 2010 Workshop on Pervasive User Modeling and Personalization (PUMP-10 at UMAP-10). (2010)
13. Bohnert, F., Zukerman, I.: Personalised pathway prediction. In: Proc. of the 18th Intl. Conf. on User Modeling, Adaptation, and Personalization (UMAP-10). (2010) 363–368
14. Bohnert, F., Schmidt, D.F., Zukerman, I.: Spatial processes for recommender systems. In: Proc. of the 21st Intl. Joint Conf. on Artificial Intelligence (IJCAI-09). (2009) 2022–2027
15. Zukerman, I., Albrecht, D.W.: Predictive statistical models for user modeling. UMUAI **11**(1-2) (2001) 5–18
16. Carmichael, D.J., Kay, J., Kummerfeld, B.: Consistent modelling of users, devices and sensors in a ubiquitous computing environment. UMUAI **15**(3-4) (2005) 197–234
17. Schmidt, D.F., Zukerman, I., Albrecht, D.W.: Assessing the impact of measurement uncertainty on user models in spatial domains. In: Proc. of the 17th Intl. Conf. on User Modeling, Adaptation, and Personalization (UMAP-09). (2009) 210–222
18. Berkovsky, S., Baldwin, T., Zukerman, I.: Aspect-based personalized text summarization. In: Proc. of the 5th Intl. Conf. on Adaptive Hypermedia and Adaptive Web-Based Systems (AH-08). (2008) 267–270
19. Fleck, M., Frid, M., Kindberg, T., O'Brien-Strain, E., Rajani, R., Spasojevic, M.: Rememberer: A tool for capturing museum visits. In: Proc. of the 4th Intl. Conf. on Ubiquitous Computing (Ubicomp-02). (2002) 379–385

The Case for Activity Models

Kurt Partridge, Oliver Brdiczka

Palo Alto Research Center, 3333 Coyote Hill Road,
Palo Alto, CA 94306, USA
{Kurt.Partridge, Oliver.Brdiczka}@parc.com

Abstract. In this position paper, we argue that user modeling systems would benefit from explicitly modeling user activities. We give examples where systems without activity inference fail to perform as desired, and briefly describe both benefits (combining data from multiple sources, learning from multiple users, improving code and technique reuse, and collecting user feedback) and costs (system complexity, development time, and inference limitations).

Keywords: User modeling, activity inference.

1 Introduction

In this position paper, we argue for including an *activity model* as a standard component in a user modeling system. An activity model provides an explicit representation of the user's behavior. Systems that infer user activity do so as an intermediate step between data collection and updating the user model. Inferring and using this explicit representation can bring several benefits to a user modeling system, which we describe in Section 3.

Activity inference is particularly important in pervasive computing user modeling systems. In these systems, a user may interact more with her environment than with the user interface of the user modeling application. With sensor data providing only indirect evidence about the user's current state, a user modeling system must interpret the data carefully to avoid making incorrect inferences.

2 Failures of Inferring User Models Directly from Data

Many user modeling systems today have no activity model. Instead, they collect information through user interactions with a system, and directly update the user model based on these interactions. The update mechanism must be determined in advance by the system designer, and in a way that is specific to the application or the domain. Although it is easier to construct systems in this way, this architecture can lead to problems. For example:

- With the greater availability of mobile location data, advertising networks are now offering "Location-Based Advertising," in which advertisements are delivered to mobile devices within a particular geographic region. For example, a coffee shop might send ads to users near their store. However, this approach presupposes that location implies an interest in the company's product. This assumption fails for the store's employees. Ad spending on them is unlikely to lead to additional revenue. But a system that explicitly modeled activity (say, by observing location traces over time to identify employment behavior) could better differentiate customers from employees and therefore deliver advertisements more effectively.

- Tolmie et al. [1] give an example of how a simple door knock can mean different things depending on the context and the expectations that have developed over time. A system draws conclusions from a simple door knock (or lack of an expected door knock) could easily misinterpret the meaning of the event. However, a correct interpretation would be possible if the system also considered the users' shared contextual history, social norms, customs, and affordances.

- In the domain of computer use (e.g., manipulation of documents, emails, web pages, or contact information), work processes and task management research such as TaskTracer, CAAD, CALO, and Activity-Based Computing have used explicit representations of work tasks and activities to improve resource prediction. Without these higher-level abstractions, grouping and association among computer use events and artifacts would likely require more connections and a more complex temporal structure, leading to less inference stability and worse detection.

3 Using Activity Models in User Modeling

We see several benefits to including activity models in user modeling systems:

- **It is easier to combine data from multiple sources.** As the number of data sources increase, updating a user model becomes increasingly difficult. Hand-constructed rules are more likely to conflict or have unforeseen exceptional conditions, and automated machine learning techniques have difficulty scaling. Inferring activity as an intermediate step can simplify the process, and research on activity recognition has addressed some of the inference challenges.

- **It is easier to combine data from multiple users.** Observations can have different meanings for different users. By developing systems that know how to map these observations into common interpretations of activities, data from one user can improve the user modeling of other users.

- **There are more opportunities for code and technique reuse.** Because activity inference can handle some of the complexity of interpreting the meaning of observations, an activity inference system can replace some of the domain-specific logic involved in user modeling.

- **It is easier to collect feedback from users.** An "activity" is a linguistic representation of what a person is doing. It has meaning to users, so users can give feedback about than a more detailed representation of the data sources that led to the activity inference. A good example for this is the CAAD system, where the user interacts with a representation of his/her activities by associating or disassociating different resources with the tasks.

Of course, the benefits of activity models are not without costs:

- **System complexity is greater.** Including an extra step between collecting observations and updating the user model does add to the system complexity. This can make debugging problems more difficult.

- **Development time may be longer.** Initially, developers will need extra time to understand how to construct systems that include activity inference. With practice, of course, this will get easier. However, important decisions like the method used for activity inference and the target set of activities require experience and some careful thought if the system is to achieve the best accuracy.

- **Some inferences may not be possible.** There is more flexibility in performing user model changes directly from observations. Requiring that updates depend only on the inferred activity can be limiting, particularly if the set of allowed activities is small. An alternative is to allow attributes to be assigned to activities, however this approach further increases the system's complexity.

Overall, we believe that in many cases the benefits will outweigh the costs. We encourage researchers and developers of user modeling systems to review recent research in activity inference, and to consider incorporating these mechanisms into their own systems.

References

1. Tolmie, P., Pycock, J., Diggins, T., MacLean, A., and Karsenty, A.: Unremarkable Computing, CHI (2002)

Adaptation Step-by-Step: Challenges for Real-time Spatial Personalization

Willem Robert van Hage[1], Natalia Stash[2], Yiwen Wang[2], and Lora Aroyo[1]

[1] VU University Amsterdam
wrvhage@few.vu.nl, l.m.aroyo@cs.vu.nl
[2] Eindhoven University of Technology
n.v.stash@tue.nl, y.wang@tue.nl

Abstract. In this paper we outline challenges for user modeling and personalization with spatial information. To illustrate those challenges we use a use case with a real-time routing system that implements a mobile museum guide for providing personalized tours tailored to the user position inside the museum and her interests. In this scenario we combine on the one hand (1) interactive discovery of user's interests applied for semantic recommendations of artworks and art-related topics, and on the other hand (2) dynamic step-by-step adaptation of a user's route through a museum based on her current position and changing interests. For this, the existing CHIP mobile museum guide was extended with a routing mechanism based on the SWI-Prolog Space package.

1 Challenges in Dynamic Spatial Adaptation

When a visitor moves around in a museum, exploring the collection of artworks based on her interest in these artworks, many spatial aspects can be considered in the process of recommending the visitor what to see. The same holds for a mobile museum guide, where we exploit techniques of dynamic personalization for such recommendations. Adding the spatial aspects allows to improve the dynamic adaptation further. For example, the assessment of the user's constantly evolving interests can be improved by step-by-step spatial information.

If we include spatial considerations into the many different elements of dynamic adaptation, we have many opportunities, but we also have as many research challenges. The opportunities we illustrate in this paper are based on our experience with the concrete museum demonstrators from our work in the CHIP project and we have elicited research challenges from that same experience.

- **Dealing with real-time information**: When we deal with real-time information in the process of dynamic adaptation, we can consider the user's position, her context, and her social interaction:
  - **User location**: Detecting a user's location inside the physical space is a first challenging step. In a museum this requires a positioning system that considers the boundaries and constraints (i.e. the walls, doors, stairs) of the space. In our case, we have an indoor space, and therefore methods

using different hardware solutions have been proposed to increase the accuracy of the indoor user positioning.
  - **User context**: For the adaptation in the collection-based semantic recommendation process, identifying the relevant context is important. This context typically includes the identification of artworks in her neighborhood, the artworks that have been already seen, the time she has already spent in the museum and additional temporal constraints (e.g. how much time is available), her general interests in art, and potentially also her physiological and emotional state [2]. The main challenge that we met in our case is to know the relevant context from a spatial perspective.
  - **Social interactions**: Social interactions can play a role in the adaptation process, e.g. the interactions between people that visit the museum together, and the spatial aspect can impact the process even more. Think of routes that the people take and potential meetings that they might have inside the museum. The social interactions are relevant both as input and output of the adaptation process and including the spatial aspect in these social interactions offers a challenging improvement.
- **Coping with the limited resources of a mobile device**: The dynamic adaptation and context-awareness ask a lot from the infrastructure and the limited resources that mobile devices offer. We have experimented in CHIP with different re-routing algorithms for the purpose of adaptation. The current algorithm can provide re-routing of a tour of artworks based on the user's position. It would be interesting however to consider more complex algorithms that would also take user preferences into account and possibly decide to add additional artworks to the tour that might be interesting for the user based on the user's close proximity.
- **Overcoming mobile platform dependance**: The personalization and recommendation process is based on knowledge about the content and the user. To reduce complexity and to ensure reusability of the knowledge representation and inference mechanisms, a flexible web-based approach is required that allows different types of systems to exchange and augment information on users and particular situations [2]. Also, the web-based architecture allows the use of multiple types of mobile devices.
- **Maintaining distributed user model**: The personalization in an application like the ones we consider here, is not a single standalone one. The user model that is relevant for this kind of dynamic adaptation in CHIP asks for the capability to exchange and integrate user model knowledge in a distributed fashion. For this reason, we have chosen a distributed and open web-based solution for user model knowledge representation.
- **Integration with third-party applications**: It would be interesting to consider technologies like Google Goggles[3] to show information about an artwork when the user points with her device to it. This is an example of an integration with third-party applications, and in our case we have chosen an approach that facilitates this space-oriented integration, that as we will show later is mainly based on offering standard interfaces and interoperability.

[3] `http://www.google.com/mobile/goggles/`

– **Evaluating in real-life settings**: In order to improve this real-time adaptation process, patterns of user's navigation and evolution of interests would be very helpful. However, collecting large volumes of such data over long periods of time is very difficult.

In the remainder of this paper, we describe how we brought the CHIP demonstrators to a next version that includes the spatial dimension.

2 Spatial Personalization in Museums: Related Work

Museum curators typically would offer tours on different topics based on the highlights of the collections. Thus, the resulting tours are characterized by a predefined selection and fixed sequence of artworks. An audio tour would still offer a predefined selection of artworks, however it allows for determining your own sequence of artworks. A number of museums, e.g. Tate Modern, Science Museum Boston, are exploring the potential of personalized museum guides. Personalized virtual tours, on the other hand, help visitors construct their own narratives[4], however they are only limited to online collections. Multimedia guides provide a promising alternative to bridge the gap between the visitor's interests and the static museum tours by using personalization techniques [5]. An adaptive mobile museum guide acts as a museum expert and provides the user with information adapted to the current situation [2]. For example, the *MIT Media Lab*[5] audio and visual narration adapts to the user's interest acquired from the physical path in the museum and length of the user stops. The mobile museum guides developed within *Hippie* [3] and *PEACH* [4] projects provide content adaptation based on technical restrictions of specific presentation devices as well as visitor's preferences and knowledge. The mobile museum guide built within *Sotto Voce* [1] project takes into account the special needs of groups visiting a museum and facilitates social interaction between group members. Another example for fostering social interaction between visitors is given by the *AgentSalon* [6] system, and *ARCHIE* [10] provides a socially-aware handheld guide that stimulates interaction between group members. The *Kubadji* mobile tour guide[6] uses a collaborative filtering approach for predicting visitor's viewing times of unseen exhibits from his viewing times at visited exhibits. The context-aware museum guide in [11] is adapting by dropping artworks if the visitor falls behind the tour or is suggesting additional artworks or taking a break at a nearby restaurant if the visitor has extra time. The environment also supports p2p interactions between visitors, to find each other, share ratings and comments about exhibits.

Important here is the fact that spatial information is used in relatively limited aspects for adaptation. Usually, the real implementation of such approaches depends on the availability of an indoor localization of people and objects.

[4] Virtual Museum (of Canada), `http://www.museevirtuel-virtualmuseum.ca/`
[5] `http://www.media.mit.edu/`
[6] `http://www.kubadji.org/`

3 Use Case: Space-CHIP Step-by-Step Route Adaptation

The CHIP project is a cross-disciplinary project, combining aspects from cultural heritage and information technologies to provide a personalized access to the museum collection both online and inside the museum[7], e.g., generating personalized museum tours, getting recommendations about interesting artworks to see, and quickly finding ways in the museum. One important aspect of the project is the use of a common distributed user model, which collects user interaction data and interprets it in terms of user's interests used further for generating recommendations and personalized tours. Additionally, we also use the e-Culture Semantic Search[8] open API to allow to find semantically related topics and artworks to include in the personalized tours. The Mobile Museum Guide allows users to access their tours created with the online Tour Wizard on their mobile devices in the museum. Details about the design and implementation of the CHIP Art Recommender, Tour Wizard and Mobile Museum Guide (ver 0.1 and ver 1.0) can be found in [9].

The CHIP Mobile Museum Guide (ver 0.1 and ver 1.0) can adapt to the user in many different ways, but mainly based on the known user preferences and availability in the museum. To relate the step-by-step adaptation also to the real physical space, spatial constraints have to be taken into account in the generation of both recommendations and museum tours. Suppose the user follows a tour of recommended artworks. If she provides a rating to an artwork that she sees, the CHIP demonstrator updates the user model and as consequence the list of recommended artworks. However, those versions of the Mobile Museum Guide do not take spatial constraints as well as information about already seen artworks into account for the adaptation. Further we show how we implement the Space-CHIP Mobile Museum Guide that includes adaptation with spatial constrains, `http://www.chip-project.org/spacechip`. The implementation is based on the SWI-Prolog Space package.

The basic approach in the Art Recommender is to recommend based on the estimated likelihood that the user will like the artworks. Even with a theme-based layout of the rooms in the Rijksmuseum, e.g., rooms for the Dutch republic or works by Rembrandt and his pupils, a set of recommended artworks can in reality be distributed over the entire museum. To improve the user experience, we therefore reorder the results of the Art Recommender to allow for an efficient walk from one artwork to the other. Such a route minimizes the walking effort, while maximizing the number of top recommendations. Also, it takes into account optional caps to the walking distance and the number of artworks. This helps the user to decide where to go in limited time.

Computing an efficient route through a museum is very similar to the *traveling salesman problem*. However, it is a significantly easier problem than the general traveling salesman problem. First, if you consider the artworks, rooms, doors, hallways, and stairs to be nodes in a connectivity graph (e.g. Fig. 1),

[7] See `http://www.chip-project.org/demo`

[8] `http://e-culture.multimedian.nl/`

then this graph is not fully connected, as there are walls and floors in the way. Second, from the way in which exhibits are created, it makes sense to view all works from a single room together. Third, floor transitions take a lot of effort. For these reasons there are only a few sensible paths through the museum.

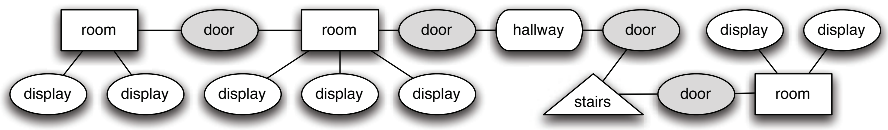

Fig. 1. Example connectivity graph

If you set the transition weight of the edges in the connectivity graph to the experienced distance instead of the actual distance, then nearest neighbor search sends the visitor to works in the same room first before making the transition to another room (or floor), which is good in the Rijksmuseum, but bad in the general case. The SWI-Prolog Space package [8] provides nearest neighbor search. As this search is unaware of the restrictions posed by the walls and floors, we base our routing on a connectivity graph search algorithm that uses intersection queries as opposed to nearest neighbor queries. First, we compute a connectivity graph between all the artworks, rooms, stairwells, etc. that considers where the doors are. Then, we compute the weighted shortest path between all artworks. The weight is based on graph distance, the type of transition (e.g. moving to another floor is more expensive), and on the distance between locations inside a room. This shortest distance matrix is used to compute an efficient path along all the recommended artworks. The exact method for route calculation is as follows:

- Pre-compute artwork distance matrix once
  1. define that stairs, hallways, toilets, are rooms
  2. define works that are on display in the museum
     (a) give the artwork a $\langle x, y, z\rangle$ coordinate
  3. define what it means to be connected
     (a) places (displays, doors) space_intersect with same room
     (b) places are stated to be connected by A chip:connectsTo B
  4. assert A chip:connectsTo B for each connected pair $\langle A, B\rangle$
  5. make connectivity graph of chip:connectsTo
  6. compute weights for each transition
     (a) graph distance plus distance within room
     (b) door transitions get a higher graph distance than artwork-artwork transitions
     (c) stairs transitions get an even higher graph distance
  7. compute and cache upper triangle matrix of weighted graph shortest path distances between all places

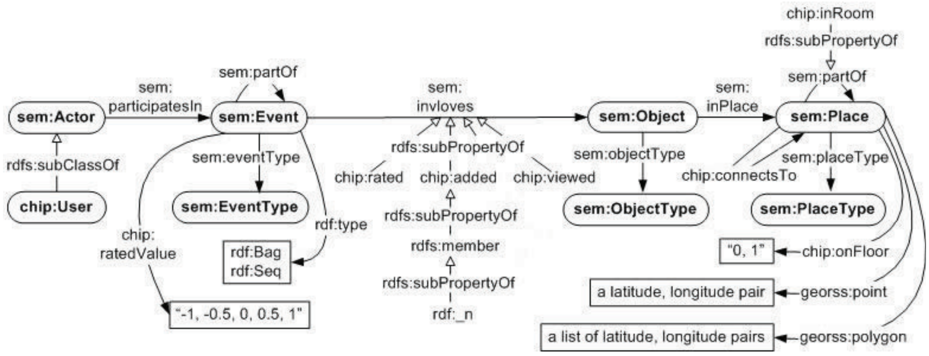

Fig. 2. Mapping CHIP user model (UM) to the simple event model (SEM)

– Apply routing algorithm for each request
  1. fetch set of recommended artworks (given by Art Recommender)
  2. fetch current position (given by user interface)
  3. fetch remaining time in museum (given by user interface)
  4. fetch maximum number of artworks to route (given by user interface)
  5. greedy nearest neighbor search in weighted distance graph until list of recommended artworks is empty:
     (a) look up the nearest recommended artwork
     (b) remove artwork from list of candidates
     (c) add path from current position to the artwork to recommended route
     (d) set current position to the location of the artwork
     (e) add length of path to total length of recommended route
  6. while total path length of recommended route takes longer than remaining time in museum
     (a) remove furthest artwork from current position
     (b) apply greedy nearest neighbor search again (step 5)

In order to provide data exchange between CHIP and the SWI-Prolog Space package we mapped (see Fig. 2) the original CHIP user model (UM) [9] to the Simple Event Model (SEM)[9] which is proposed by van Hage et al. [7].

For example, when the user rates with four stars both the painting "Woman Reading a Letter" and its creator Johannes Vermeer, this results in a list of recommended artworks, which further used to generate the Tour of Recommended Artworks (see Fig. 3). We use icons in a different color to indicate artworks that are in the tour and connect them with the tour line. The user location is indicated with an icon at the entrance door on the ground floor. During the visit the user views artworks that are in the tour but is also attracted by other artworks outside her tour. In that case, the tour may be re-routed taking into account the user's interest in these additional artworks. Similarly, the user can also rate any

[9] For this work we use this version: `http://semanticweb.cs.vu.nl/2009/04/event/`. A newer version is available at `http://semanticweb.cs.vu.nl/2009/11/sem/`.

artwork she sees on her way. These actions result in the tour being dynamically adapted taking into account the history of seen artworks, changing interests and current location. Thus, if the user likes the works by Frans Hals and Ferdinand Bol that she comes across on her way to the recommended Johannes Vermeer works, she can submit new rating and this automatically updates the tour. The updated tour is shown in Fig. 4. For the sake of clarity we have highlighted the works from the original tour with red, the new Frans Hals recommendations with yellow and the new Ferdinand Bol work with blue.

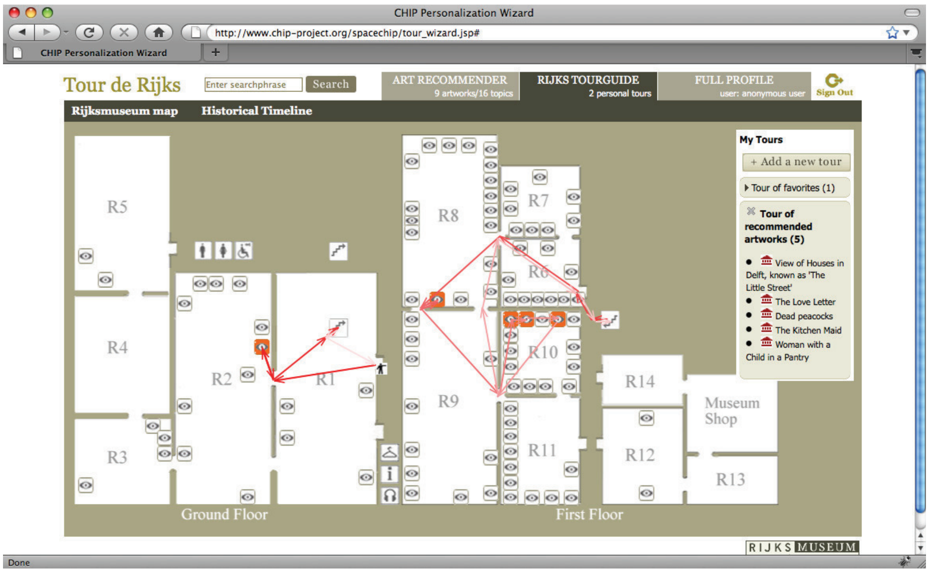

Fig. 3. Initial route of the tour of recommended recommended artworks

In [12] we evaluated (1) the usefulness of recommendations to the users and (2) the efficiency of the route calculation.

Acknowledgments. The work on the CHIP project is funded by the Dutch Science Foundation CATCH program and is in collaboration with the Rijksmuseum Amsterdam. Part of this work has been carried out as a part of the Poseidon project in collaboration with Thales Nederland, under the responsibilities of the Embedded Systems Institute (ESI). This project is partially supported by the Dutch Ministry of Economic Affairs under the BSIK program.

References

1. Aoki, P.M., Grinter, R.E., Hurst, A., Szymanski, M.H., Thornton, J.D., Woodru, A.: Sottovoce: exploring the interplay of conversation and mobile audio spaces. In: SIGCHI Conference on Human factors in computing systems (2002)

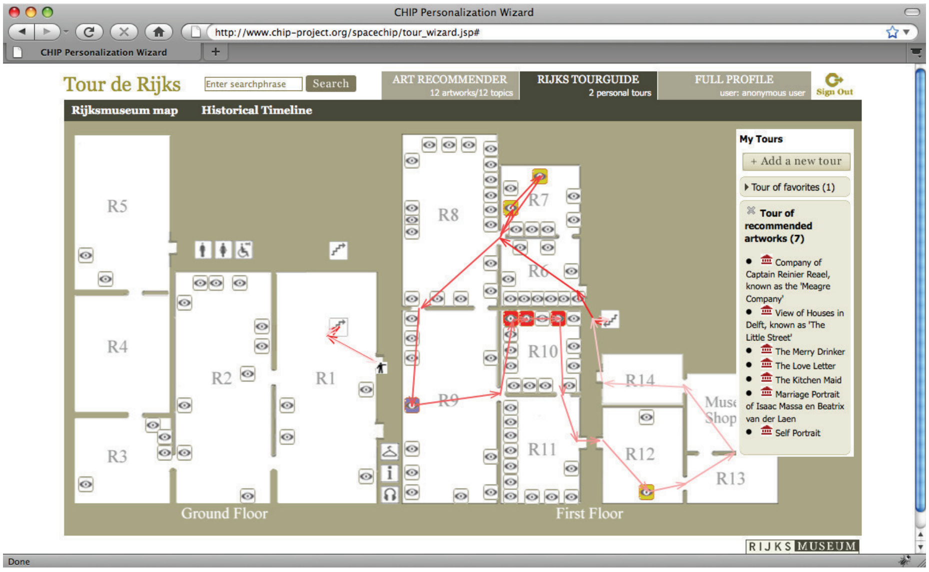

Fig. 4. Re-routed tour of recommended artworks

2. Kruger, A., Jorg, B., Heckmann, D., Kruppa, M., Wasinger, R.: Adaptive Mobile Guides. In: The Adaptive Web, pp. 521–549. Springer (2007)
3. Oppermann, R., Specht, M.: A nomadic information system for adaptive exhibition guidance. In: ICHIM99 Conference (1999)
4. Rocchi, C., Stock, O., Zancanaro, M., Kruppa, M., Kruger, A.: The museum visit: generating seamless personalized presentations on multiple devices. In: IUI (2004)
5. Roes, I., Stash, N., Wang, Y., Aroyo, L. : A personalized walk through the museum: The CHIP interactive tour guide. In: ACM CHI-Student Competition (2009)
6. Sumi, Y., Mase, K.: Interface Agents that facilitate knowledge interactions between community members. In: Cognitive Technologies Series, pp. 405–427 (2004)
7. van Hage, W.R., Malaisé, V., de Vries, G., Schreiber, G., van Someren, M.: Combining ship trajectories and semantics with the simple event model (sem). In: 1st ACM International Workshop on Events in Multimedia. Sheridan Publishers (2009)
8. van Hage, W.R., Wielemaker, J., Schreiber, G.: The space package: Tight integration of space and semantics. In: 8th ISWC Workshop: TerraCognita (2009)
9. Wang, Y., Stash, N., Aroyo, L., Gorgels, P., Rutledge, L., Schreiber, G.: Recommendations based on semantically enriched museum collections. Journal of Web Semantics. 6(4), 43–58 (2008)
10. Luyten, K., Van Loon, H., Teunkens, D., Gabrils, K., Coninx, K., Manshoven, E.: ARCHIE: Disclosing a Museum by a Socially-aware Mobile Guide. In: 7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (2006)
11. Chou, S.-C., Hsieh, W.-T., Gandon, F., L., Sadeh, N.M.: Semantic Web Technologies for Context-Aware Museum Tour Guide Applications. In: 19th Int. Conference on Advanced Information Networking and Applications, 709–714 (2005)
12. van Hage, W.R., Stash, N., Wang, Y., Aroyo, L.: Finding Your Way through the Rijksmuseum with an Adaptive Mobile Museum Guide. In: ESWC (2010)

Towards Life-long Personalization Across Multiple Devices: The Case of Personal Career Management

Rainer Wasinger[1], Anthony Collins[1], Michael Fry[1], Judy Kay[1], Tsvi Kuflik[2], Robert Kummerfeld[1]
[1] CHAI Group, School of IT, Sydney University, 2006, NSW, Australia
{rainer.wasinger, anthony.collins, michael.fry, judy.kay, bob.kummerfeld}@sydney.edu.au
[2] University of Haifa, Haifa, Israel
tsvikak@is.haifa.ac.il

1 Illustrative Scenario

Consider the following scenario. Sammy is a professional software developer, and although content with his current job, is also keen to keep his skills up to date and to keep a look out for new career opportunities in both his current field of work (i.e. as a software developer) as well as potential new fields of work (e.g. as either a developer specialising in user-models and mobile applications; or as a project manager).

Sammy has undergone many years of education, having completed 6 years of primary school, another 6 years of high school, 4 years at a University, and a number of speciality post-graduate short-term courses. He has much of these 16+ years of education (as well as details on his past job positions and his possessed skills) condensed into a 2 page curriculum-vitae. Sammy has an inkling that having 16+ years of education and experience, condensed into 2 pages of curriculum-vitae might be missing something, but he can’t quite put his finger on what (or just how much).

Now consider the following vision: Sammy has a life-long user model containing details of his past education and employment, and at a granularity that contains individual course names and all the topics covered within each course, according to a given curriculum (be that a national school curriculum, a university-specific curriculum, or other), as well as details of past work experiences (types of systems, programming languages etc). He has just registered to an online career web service and installed a career management application onto his favourite mobile device and is keen to look at what career prospects might exist for him both now and in 5 years time. Using data contained within his life-long user model, the system is able to quickly configure an application-specific persona that is used to filter thousands of job listings from a 3rd-party server. He then commences to browse these personalised listings on his mobile device whenever he has a few spare minutes of time (e.g. during his bus trip to and from work) and notes to himself the importance of being able to efficiently consume information on small screen devices and in busy mobile contexts, as well as the ability to easily pause and resume tasks when the user is mobile. He also appreciates how a number of different and diverse data sets (e.g. his own life-long user model and the 3rd-party job listings) are being used to create him a personalised list of easily browsed career opportunities within his organisation, the city he lives in, as well as more globally; and the way that educational opportunities also link back to current weaknesses in his career management profile.

2 Discussion

The above described scenario is not too fictitious. It demonstrates one of the aspects where lifelong user modelling is needed for supporting a personalized service for a user on a daily basis. There are many similar activities in a user's life that will benefit from being able to harness information contained within a life-long user model; career management - which we feel has strong ties to education as well as work experience and other contextual aspects - is one such task. The research community is still a long way off from solving all the issues surrounding the use of life-long user models. [1] describes some of these issues as being: data capture, user profiling, reasoning, and recommendation, and [2] describes the issues of interoperability, scrutability, control, and user privacy. This position paper presents some of the technical and social implications and possible solutions for integrating multiple data sources with a life-long user model, to be used with small-screen devices in pervasive computing. Following are some suggested solutions for challenges implicitly suggested by the above scenario.

From a **user modelling** point of view we need to consider what user-modelling data will look like for educational and career-management domains and how it may be maintained over time. As representations and terminology tends to evolve, simple evidence-based user models will require considerable reasoning efforts for combining relevant pieces (as well as considerable storage). One possibility may be to rely on domain ontologies that are flourishing nowadays and represent users' skills as an overlay over domain ontologies, where the user model contains a link to the ontology as well as a list of abstract terms and a definition of levels of knowledge and/or interests. This option solves the need for detailed domain ontologies or relying on reasoning about personal definitions. Still, reasoning will be needed for integrating data from different ontologies (different institutions will probably use different ontologies), but this will be largely resolved by future ontology matching research. Periodic maintenance will take care of ontologies evolvement over time.

Place of storage of user modelling data is another challenge. Users should have control over their personal information, hence the lifelong user model should reside in the possession of the user (e.g. on the personal device and/or on a personal computer or on a secure server). This approach, combined with the previous suggestion relies heavily on internet access (overlay model over remote ontologies), however, nowadays it is reasonable to assume that ongoing communication is possible whether by using a mobile service provider or WiFi. When considering the home environment, profiles may reside on the home-server and be synchronized with the mobile version every time the user enters/leaves the house. Having a personal profile stored over a secure server may be a possibility as well. The key concept is that the lifelong user model will be in the possession of the user and information will be revealed to service provider following user-defined privacy policies.

For a life-long user model to be useful, relevant communication protocols for information sharing are essential. An **API** needs to exist through which applications can contribute and retrieve data. Additional to such an API, a supporting framework to control access to (and to monitor and record) the data that is and has been provided over time to individual applications should also exist. Such communication can be based on extensions to the already suggested UserML.

In addition to the actual user model data, it is conceivable that **reasoning and inference** layers will exist on top of the data and rules be applied to the data, in order to increase the relevance of that data to the user's current context. For example, consider the concept of "forgetting" learned material and how this might apply to both technical skills and social skills that a user has acquired over time. The effect of forgetting might for example be modelled as a rule based on elapsed time, though perhaps also using information in a forgetfulness layer so that each user model attribute can be treated uniquely. The notion of privacy and access control may also be modelled using a layer approach in which users indicate by privacy policies, which attributes can be used for which purposes (e.g. not to be used at all, only for intermediary purposes, or only by certain applications).

Visualisation of user-modelling data and user interface design will be particularly important considering for example, on top of the limited display space provided by current mobile devices, that the user model data is likely to be quite exhaustive, and viewed at different levels of granularity (e.g. consider sensor data), and on different time scales (e.g. data now and as it was in the past), and may (depending on the intelligence intertwined in the user model) be a result of multiple intermediary calculations on other attribute values. User alerts and user explanations will also need to be carefully crafted, e.g. to explain the type of data, the quantity of data, past requested-data, and expected future data requests by 3rd-party applications. Similarly, the formulation of reasons as to why an application requires a particular user model attribute, and consequences (for the application) in the user denying one or more of the attribute values would need careful consideration, as too features for a supporting framework to help prevent a user (e.g. in the form of user alerts or data locks) from providing 3rd-party applications with too much information (or even their complete user model) at the expense of a short-term gain (e.g. "give me all your data and I'll let you download this file now"). One solution for dealing with these UI design issues and interaction flow sequences will be to conduct careful user evaluations to determine what works best for users of small-screen devices in pervasive computing settings.

We discussed a wide range of challenges posed by applying lifelong user modelling in practice, in a pervasive scenario and suggested what seems to be initial feasible scenarios, however, we are sure that there are additional challenges, as well as solutions that we look forward to discuss.

References

1. Elliott, D., Hopfgartner, F., Leelanupab, T., Moshfeghi, Y., Jose, J., An Architecture for Life-long User Modelling, In: Proceedings of the Lifelong User Modelling Workshop at User Modelling, Adaptation, and Personalisation, pp. 9 – 17 (2009).
2. Kay, J., Lifelong Learner Modelling for Lifelong Personalized Pervasive Learning. In: IEEE Transactions on Learning Technologies, 1(4), pp. 215 – 228 (2008)

Author Index