

The CHES Project: Adaptive Personalized Storytelling Experiences in Museums

Maria Vayanou¹, Manos Karvounis¹, Akrivi Katifori¹, Marialena Kyriakidi¹,
Maria Roussou¹, Yannis Ioannidis¹

¹ MaDgIK Lab, Department of Informatics and Telecommunications
University of Athens, Panepistimioupolis, Ilissia, 15784 Athens, Greece
{vayanou,manosk,vivi,marilou,mroussou,yannis}@di.uoa.gr

Abstract. In this work, we describe the basic elements of an effort towards achieving personalized storytelling for museum visits in the context of the CHES project, with a focus on the profiling techniques employed.

Keywords: Profiling, Adaptive Storytelling, User Modeling, Personalization

1 Introduction

CHES (Cultural Heritage Experiences through Socio-personal interactions and Storytelling) is a research prototype that was developed under the CHES project (<http://www.chesexperience.eu/>), aiming to enrich museum visits through personalized interactive *storytelling*. It uses a) personalized information about cultural artefacts to create customized stories that guide individuals or groups through a museum and b) aspires to (re-)inject the sense of discovery and wonder in the visitor experience. The CHES system employs mixed reality and pervasive games techniques, ranging from narrations to augmented reality on smart phones. Two museums participated in the effort, each with a different scope and end user requirements: the Acropolis Museum in Greece, and the Cité de l'Espace in France.

There are two types of CHES users, namely the visitors, who “consume” CHES experience through their web or mobile terminals, and the authors, who design the experiences. CHES is following a hybrid, plot-based approach with pre-defined content, where *story authors* (curators, museum staff, script writers) write stories around pre-selected museum themes.

2 Authoring CHES Stories and Storytelling Model

Similarly to the making of a movie, the creation of CHES stories includes four main phases, namely scripting, staging, producing and editing. During scripting, the author chooses the main story concepts, sketches the plot and writes the narrative text, i.e., the script. In staging, the author associates parts of the script with exhibits, paths and other spots in the museum environment. Then, a set of multimedia resources are

produced for the staged script, including audiovisual materials, games, quizzes, augmented reality applications, referred to as activities. Finally, the author does the montage, selecting and ordering the activities to realize the script.

In correspondence to the authoring phases, stories are represented as graphs in three different levels of abstraction, namely, the scripting, the staging, and the editing graphs, defining the succession of their atomic pieces and enabling conditional branching based on a variety of events or/and visitor characteristics, over all the three levels. The three graphs are interlinked, so the combined graph forms the story's Storytelling Graph. The overall CHESS Storytelling Graph (CSG) starts with a branching point which leads to all the CHESS stories authored so far.

All the CSG entities (i.e. their atomic pieces as well as the graph branches) are annotated with author selected features. Several features have been exploited so far, such as the topic, information type (real or fictional), script tone, connection to exhibits, required user role, duration, multimedia type, etc.

3 Personalized and Adaptive CHESS Visitor Experience

A typical CHESS experience starts as soon as the visitor enters the museum environment. The visitor goes to a specific web location with his tablet where he is required to log-in into the CHESS application and fill out a short quiz, to gather initial evidences regarding his/her preferences. Then the Adaptive Storytelling Engine (ASTE) starts traversing the CSG graph. Whenever a branching point is met, the ASTE performs two main steps: i) evaluates any hard constraints expressed on each branch to identify the valid ones (e.g. a branch may be unavailable for children) and ii) estimates the visitor's interest in the valid branches to rank them accordingly.

Aiming to reach the right balance between the mental load created to the visitor by the presentation of numerous questions and fully automated decision making, branching points are annotated by the authors as *mandatory*, *automatic* or *optional*. When a mandatory branching point is reached, the ASTE generates a menu where the available options are ranked according to the visitor's profile, while highlighting the first one. In automatic ones, the ASTE makes a decision without informing the visitor about the available options. Finally, when an optional one is reached, the ASTE decides whether a menu will be displayed or an automatic decision will take place. An automatic decision is taken if (a) there exists one option that is significantly better than the rest ones for the current visitor, and choosing it will not omit other story parts that the visitor may also like, or, (b) there exists only one option that the visitor will most probably like and the rest of the options will most probably be disliked.

Depending on the visitor's choice or the ASTE's decision, the CSG is traversed accordingly and the appropriate multimedia resources are fetched and presented to the visitor's terminal.

4 Matching CHESS Visitors to CHESS Content

The visitor's profile contains information about past actions and demographic data, as well as his preferences over the objects he has interacted with (i.e. the CSG entities),

referred to as *Interaction Objects*. Visitor's actions are interpreted and a preference value in $[-1,1]$ is extracted. To estimate the visitor's interest in a set of *Candidate Objects*, we use the well-known k-nearest neighbor recommendation algorithm [1]. First, we calculate the similarity of each Candidate Object to each Interaction Object in the visitor's profile and for each Candidate Object we keep the k Interaction Objects with the largest similarities (we have used cosine similarity though any proper similarity metric may be employed). Given the Interaction Objects with the k largest similarities, we calculate the predicted preference value for the corresponding Candidate Object utilizing a weighted average.

In this way, visitor profiling and matching to CSG entities is not closely tied to the actual features' values used by the authors for annotation. Authors are enabled to have an open tag vocabulary, which they can specify upon their understanding of the current story; the only requirement is to use the same vocabulary throughout the story.

To achieve this we have used PAROS [2], a system that builds and maintains user profiles following a generic, graph-based user modeling framework.

5 Profile Initialization and Story Selection

To support the authoring process and address the personalization cold start problem, CHESS utilizes the notion of *personas*, a design tool from the marketing world. The system leverages persona definition to match visitors to personas, essentially aligning visitor preferences to the author's understanding of the museum visitors [3]. This approach has been applied in an evaluation study that took place in the Cité de l'Espace and it is described in [4]. We have also explored an alternative approach, interpreting visitors' answers to the initial quiz as evidences about their likes in fictional story parts, which are appropriately annotated. Initial story selection is then conducted by matching the visitor's profile to the corresponding stories' annotations. Results from a recent evaluation study with 24 participants in the Acropolis Museum showed that the adopted approach reached approximately 82% of correct decisions.

6 Profile Update under the CHESS Experience

The CHESS Profiler monitors the visitor's behaviour, interprets it as negative or positive feedback, adjusts the visitor's profile accordingly and uses the updated profile in the rest of the experience. The following visitor actions are exploited on that front: skipping (interpreted as high negative feedback on the corresponding activity), completion of activity (low positive feedback on the activity), menu selections (high positive feedback on the script branch) and non-selections in menus (low negative feedback on the corresponding script branches). Based on the visitor's actions on all the activities that comprise a script unit, the Profiler estimates the visitor's preference on the current script unit. Moving a step further, the script unit preferences are used together to infer the visitor's preferences on the whole script branch.

However, due to the big amount and diversity of entities included in the visitor's experience, which may last from 15 minutes to 1 hour, implicit feedback on its own may lead to inaccurate conclusions. For instance, evaluation results have shown that

skipping actions may actually occur due to the visitors' dislike in previous parts of the story, rather than on the ones that were actually skipped. Accounting for the main issues observed during the evaluation studies, we have also implemented a conditional explicit feedback approach. A feedback dialogue is dynamically injected into the user's experience through menus, when certain conditions are met, aiming to increase profiling accuracy.

For instance, if the visitor skips many activities in a short time then a menu is shown asking if he disliked the story part, if he is getting tired and would like to shorten the experience, or if he already knew the upcoming story part. Similarly, if the visitor skips a story part that his/her profile indicates he will like, a menu is shown asking if he is getting tired, if he disliked only the parts of the story he skipped but liked the rest, or if he disliked the whole story part. The main strength of this approach is that explicit feedback is requested only when needed, thus minimizing story interrupts and feedback overhead, while feedback requests occur as the system's response to the visitor's actions.

To further increase the accuracy of implicit feedback, we have also leveraged the visitor's location and movements, the time he spends at each story part, and the way he holds the tablet, examining the visitor's viewing angle. Due to the lack of precise location tracking these techniques have been implemented as a proof of concept and they are showcased with a Javascript demonstration running in the tablet.

7 Conclusions

The provision of personalized content for storytelling experiences in museums entails several profiling challenges. Explicit feedback needs to be maintained minimal to avoid fragmenting the story's plot. At the same time, visitor actions are the result of a guided, complicated interaction with the story's content, the terminal, the museum's space and exhibits, thus requiring for sophisticated and precise visitor monitoring techniques to detect visitor divergence and increase the accuracy of implicit profiling.

References

1. Ricci, F., Rokach, L., Shapira, B., Kantor, P.: Recommender systems handbook, Springer 2011, isbn 978-0-387-85819-7
2. Ioannidis, Y.E., Vayanou, M., Georgiou, T., Iatropoulou, K., Karvounis, M., Katifori, V., Kyriakidi, M., Manola, N. Mouzakidis, A., Stamatogiannakis, L., Triantafyllidi, M.L.: Profiling Attitudes for Personalized Information Provision. *IEEE Data Eng. Bull.* 34(2): 35-40 (2011)
3. Roussou, M., Katifori, A., Pujol, L., Vayanou, M., Egglestone, S.R: A life of their own: museum visitor personas penetrating the design lifecycle of a mobile experience. *CHI Extended Abstracts 2013*: 547-552
4. Vayanou, M., Karvounis, M., Kyriakidi, M., Katifori, A., Manola, N., Roussou, M., & Ioannidis, Y.: Towards Personalized Storytelling for Museum Visits. *6th Int Workshop on Personalized Access, Profile Management, and Context Awareness in Databases* (2012)