Engineering multiuser museum interactives for shared cultural experiences

Roberto Confalonieri [a,*], Matthew Yee-King [b], Katina Hazelden [b], Mark d'Inverno [b], Dave de Jonge [a], Nardine Osman [a], Carles Sierra [a], Leila Agmoud [c], Henri Prade [c]

[a] Artificial Intelligence Research Institute (IIIA-CSIC), Bellaterra, Spain
[b] Goldsmiths College, University of London, London, UK
[c] Institute de Recherche en Informatique de Toulouse (IRIT), Université Paul Sabatier, Toulouse, France

ARTICLE INFO

ABSTRACT

Multiuser museum interactives are computer systems installed in museums or galleries which allow several visitors to interact together with digital representations of artefacts and information from the museum's collection. In this paper, we describe WeCurate, a socio-technical system that supports co-browsing across multiple devices and enables groups of users to collaboratively curate a collection of images, through negotiation, collective decision making and voting. The engineering of such a system is challenging since it requires to address several problems such as: distributed workflow control, collective decision making and multiuser synchronous interactions. The system uses a peer-to-peer Electronic Institution (EI) to manage and execute a distributed curation workflow and models community inter-actions into scenes, where users engage in different social activities. Social interactions are enacted by intelligent agents that interface the users participating in the curation workflow with the EI infra-structure. The multiagent system supports collective decision making, representing the actions of the users within the EI, where the agents advocate and support the desires of their users e.g. aggregating opinions for deciding which images are interesting enough to be discussed, and proposing interactions and resolutions between disagreeing group members. Throughout the paper, we describe the enabling technologies of WeCurate, the peer-to-peer EI infrastructure, the agent collective decision making cap-abilities and the multi-modal interface. We present a system evaluation based on data collected from cultural exhibitions in which WeCurate was used as supporting multiuser interactive.

© 2015 Elsevier Ltd. All rights reserved.

1. Introduction

In recent times, high tech museum interactives have become ubiquitous in major institutions. Typical examples include augmented reality systems, multitouch table tops and virtual reality tours (Gaitatzes and Roussou, 2002; Hornecker, 2008; Wojciechowski et al., 2004). Whilst multiuser systems have begun to appear, e.g. a 10 user quiz game in the Tate Modern, the majority of these museum interactives do not perhaps facilitate the socio-cultural experience of visiting a museum with friends, as they are often being designed for a single user. The need to support multiuser interaction and social participation is a desirable feature for shifting the focus from content delivery to social construction (Walker, 2008) and for the development of a cultural capital (Hope et al., 2009).

At this point, we should note that mediating and reporting the actions of several 'agents' to provide a meaningful and satisfying sociocultural experience for all is challenging (Heath et al., 2005). Social interaction and collaboration are key features for the development of a socio-technical system like the one described in this paper. On the one hand, the system has to enhance user interactions and should be accessible independently from user locations. This requires a robust and flexible infrastructure that is able to capture a *social workflow* and the *dynamics* of the community which will engage in the system. On the other hand, the system has to assist users in *collective decision making* and *negotiation*, and to foster participation and discussions about the cultural artefacts. This requires the use of autonomic agents that can advocate and support the desires of their users e.g. aggregating

* Corresponding author.
E-mail addresses: confalonieri@iiia.csic.es (R. Confalonieri),
m.yee-king@gold.ac.uk (M. Yee-King), kat9@me.com (K. Hazelden),
dinverno@gold.ac.uk (M. d'Inverno), davedejonge@iiia.csic.es (D. de Jonge),
nardine@iiia.csic.es (N. Osman), sierra@iiia.csic.es (C. Sierra),
agmoud@irit.fr (L. Agmoud), prade@irit.fr (H. Prade).

opinions for deciding which images are interesting enough to be discussed, and proposing interactions and resolutions between disagreeing group members.

Another trend in museum curation is the idea of community curation, where a community discourse is built up around the artefacts, to provide different perspectives and insights (Turner, 2011). This trend is not typically represented in the design of museum interactives, where information-*browsing*, and not information-*generation* is the focus. However, museums are engaging with the idea of crowdsourcing, with projects such as "Your Paintings Tagger" and "The Art Of Video Games" (Greg, 2011; Barron, 2012), and folksonomies with projects such as "steve.project" and "Artlinks" (Hellin-Hobbs, 2010; Cosley and Lewenstein, 2008; Cosley and Baxter, 2009). Again, controlling the workflow within a group to engender discussion and engagement with the artefacts is challenging, especially when the users are casual ones as in a museum context.

In this paper, we describe WeCurate, a first of its kind multiuser museum interactive. WeCurate uses a multiagent system to support community interactions and decision making, and a peer-to-peer Electronic Institution (EI) (de Jonge et al., 2013) to execute and control the community workflow. Our aim is not only to make use of agent technology and Electronic Institutions as a means to implement a multiuser museum interactive, but also to relate agent theory to practice in order to create a socio-technical system to support an online multiuser experience.

To this end, we specify a community curation session in terms of the scenes of an EI for controlling community interactions. We support system and user decisions by means of personal assistant agents equipped with different decision making capabilities. We make use of a multimodal user interface which directly represents users as agents in the scenes of the underlying EI and which is designed to engage casual users in a social discourse around museum artefacts by chat and tag activity. We present the evaluation of the system for determining the level of interactions and social awareness perceived by the social groups while using the system, and for understanding whether our agent-based decision models can predict what images users like from their behaviour. We validate our scene-based design and, consequently, our EI model, from the social behaviour of users that emerged naturally during the curation task.

This paper unifies and develops the content of the conference papers (Amgoud et al., 2012; Yee-King et al., 2013; Hazelden et al., 2013) by describing the underlying peer-to-peer EI infrastructure and presenting an analysis of the decision making models employed by the agents. The evaluation is based on data collected from cultural exhibitions in which WeCurate was used as a supporting multiuser museum interactive. The rest of the paper is organised as follows. Section 2 provides an overview of the system, whereas Section 3, 4, 5, and 6 respectively describe the EI infrastructure and workflow, the personal assistant agents, the interface and the adopted technologies. Section 7 presents the evaluation of our system. After discussing the evaluation's results (Section 8), Section 9 presents several works that relate to ours from different perspectives. Finally, in Section 10 we draw some conclusions and we envision some of the ideas we have in mind to improve the current system.

2. System overview

WeCurate is a museum interactive which provides a multiuser curation workflow where the aim is for the users to synchronously view and discuss a selection of images, finally choosing a subset of these images that the group would like to add to their group collection. In the process of curating this collection, the users are encouraged to develop a discourse about the images in the form of

weighted tags and comments, as well as a process of bilateral argumentation. Further insight into user preferences and behaviours is gained from data about specific user actions such as image zooming and general activity levels.

A multiuser interactive is a typical example of a system in which human and software agents can enter and leave the system and behave according to the *norms* that are appropriate for that specific society. For instance, it can be desirable to have only a certain number of users taking part to a curation session or to allow each user to express at most one vote. A convenient way to coordinate the social interactions of agent communities is by means of an Electronic Institution (EI) (Arcos et al., 2005).

An EI makes it possible to develop programs according to a new paradigm, in which the tasks are executed by independent agents, that are not specifically designed for the given program and that cannot be blindly trusted. An EI is responsible for making sure that the agents behave according to the norms that are necessary for the application. To this end, the actions that agents can perform in an EI are represented as messages and are specified according to an interaction protocol for each scene. The EI checks for each message whether it is valid in the current state of the protocol, and, if not, prevents it from being delivered to the other agents participating in the EI. In this way, the behaviour of non-benevolent agents can be controlled.[1] Therefore, the EI paradigm allows a flexible and dynamic infrastructure, in which agents can interact in an autonomous way within the norms of the cultural institution.

EIs have usually been considered as centralised systems (Noriega, 1997; Esteva, 2003). Nevertheless, the growing need to incorporate organisational abstractions into distributed computing systems (d'Inverno et al., 2012) requires a new form of EIs.

In WeCurate, since users can be physically in different places, it is desirable to run an EI in a distributed manner to characterise human social communities in a more natural manner. To this end, we implemented a new form of EI that runs in a distributed way, over a peer-to-peer network (de Jonge et al., 2013). The multiuser curation workflow has been modeled as scenes of an EI and scene protocols. The workflow is managed and executed by a peer-to-peer EI, with agents operating within it to represent the activities of the users and to provide other services. The users interact with the system using an animated user interface. An overview of the system architecture, showing the peer-to-peer EI, the User Assistant agents and user interface components are provided in Fig. 1.

In the following sections, we present the internal structure of the peer-to-peer Electronic Institution and the WeCurate curation workflow. Then, we describe the agents that participate in the workflow, with particular emphasis on user representation and collective decision making. The user interface is presented with images of the different scenes in the workflow. The system architecture is described, including the connections between EI, agents and UI. Finally, the adopted technologies used to implement the system are briefly explained.

---

[1] The EI cannot control, however, the behaviour of a non-benevolent agent when it fails to perform an action that the protocol requires it to perform. It essentially cannot force an agent to do something it does not wish to do. This is because EIs are designed for autonomous agents, and although we would like agents to behave in certain ways, their autonomy must be maintained. In such a case, either the protocol engineer can make use of timeouts to make the protocols resilient against such scenarios, or misbehaviour should be addressed through other measures, such as sanctions and rewards (Modgil et al., 2009; Gaertner et al., 2007), trust and reputation (Osman et al., 2014), and so on.

The EI also cannot control the behaviour of a non-benevolent agent that *does* follow a protocol but does it in a malicious way, for instance, by pretending to like an image, or by pushing other users to change their opinion with no specific reason, etc. To address this situation, again trust models can be used to detect and block the malicious behaviour of an agent, for instance, by assessing the trustworthiness of an agent through learning from similar past experiences (Osman et al., 2014).