k:info: An Architecture for Smart Billboards for Informal Public Spaces

Max Van Kleek
MIT Computer Science and Artificial Intelligence Laboratory (CSAIL)
200 Technology Square
Cambridge MA, 02139 USA
max@mit.edu

ABSTRACT

High-traffic public spaces in the workplace are rich breeding grounds for informal collaborations among knowledge workers; yet, very little technology currently inhabits these spaces today. *k:info* is a context-aware information billboard that aims to inspire informal interactions in these spaces by providing a dynamic display of items that are relevant and easily visible to users nearby.

1. INTRODUCTION

Public gathering spaces such as lounges, elevator lobbies, and hallways are places where informal social encounters occur most frequently in the workplace[4]. In addition to serving as the crossroads for day-to-day activities, these spaces harbor a relaxed social atmosphere, where people feel naturally inclined to gather and talk casually about anything that may be on their minds. As a result, these spaces encourage social connections to be made, shared interests to be discovered, and, perhaps most importantly, informal collaborations to form among people who may otherwise never have realized the opportunity to work together. Despite the importance of such social encounters and informal collaborations in knowledge-driven organizations [1], these spaces still largely lack any information infrastructure. This inspired the Ki/o project [5] to design such an information infrastructure, which consists of an *intelligent kiosk* platform and software architecture to be integrated into these spaces.

One of the first applications being developed for the Ki/o platform is a "smart" information bulletin/billboard called k:info that opportunistically uses available contextual clues from the environment to schedule items for display. Like the advertising billboards and dynamic newspaper information displays in Stephen Spielberg's science fiction thriller *Minority Report*, this information billboard dynamically adapts its display for its audience and based on contextual changes in the environment. But unlike the *Minority Report* advertising billboards which aim to persuade, the aim of k:info billboards is to spark informal social conversations among passers-by, by displaying information that coincides with their common interests.

2. DESIGN CHALLENGES

2.1 Content selection

Realizing content personalization on billboards and other wide-audience public information displays is challenging for a number of reasons. Most contemporary personalization systems exclusively rely on statistical collaborative filtering algorithms to choose what to display. These algorithms work by logically clustering people based upon how similarly people like or dislike items they have previously seen. Statistical collaborative filtering systems usually require users to state this information explicitly, such as by having users assign scores or specify rankings to each item. Obtaining scores in such a manner, however, is impractical for high-traffic public displays, due to the reason that the interaction duration between the user and such a display is typically extremely short. Furthermore, statistical collaborative filtering algorithms suffer from the "cold-start", or "ramp-up" problem, meaning that they require a large amount of initial data about each user before they can make recommendations. The problem is compounded by the potentially unbounded size of the user-base of such public displays, as well as the large probability that any given user may not have ever used the system before. Finally, making a collaborative filtering engine situationally or environmentally *context-aware* (as defined by [6]), requires an exponential amount of data to train the system, because new users must be classified along each of the contextual dimensions.

These observations have led us to a new approach that combines a conventional collaborative filtering engine with a symbolic knowledge-based recommendation architecture. This architecture explicitly represents various states of the world, such as user profiles, and maintains heuristics that can make context-sensitive recommendations based on this knowledge.

2.2 Information Composition and Display

Content presentation, or the way information is conveyed to users, may be made context-aware as well. Information pertaining to the capabilities and characteristics of the physical kiosk display, as well as user presence information, such as how closely users are standing to the display, or how many users are nearby, can be used to optimize an article for readability. If the display is small relative to the distance users are standing from the display, for example, k:info should choose a display technique that is well-suited and readable for users at their standing distance, such as a rapid-serial visual presentation (i.e., slide-show) method, instead of a bulletin-board, collage-style layout.[2] Similarly, if the system has identified a user with a disability in the vicinity, such as a user with a visual impairment, alternate channels such as text-to-speech may be activated.

3. APPROACH

3.1 Knowledge-Based Selection

The k:info system performs two functions: *collection* followed by *selection*. Specifically, k:info must collect updated display candidates (e.g., news articles and event announcements) and choose from among them what to display at any particular moment. To make this judgment, the system also needs to collect the contextual information it requires for determining the relevance of each item to the current context. This includes information directly perceived from the physical and digital "surroundings", such as the time of day, weather outside, or user presence and identity information, as well as more static information that can be explicitly updated by the system's maintainer, such as display device characteristics and configuration. Thus, k:info requires perceptual capability, as well as a facility that allows knowledge to be inspected and updated quickly and easily.

3.2 k:info Blackboard Architecture

Knowledge-based selection requires the ability to consolidate a large assortment of heterogeneous types of information. Blackboard architectures, as popularized by the *Hearsay II* first speech recognition system are well-suited for this task[3]. Blackboards consist of independent modules, called *Knowledge Sources* (KSes) that either embody a type of expertise or represent an external data source, and which communicate new information across the blackboard, a persistent knowledge repository. A simple Java blackboard architecture called the *Context Keeper* was designed for k:info.[5]

3.3 k:info Knowledge Sources

Knowledge sources in k:info are divided into three functional categories:

1. *Perceptual* Knowledge Sources. The first, perceptual KS agents, add the lowest-level information to the blackboard. This includes data gathered from hosts on the Internet, such as news feeds (currently CNN and BBC), announcement lists (currently MIT and CSAIL Events Calendars), e-mail messages, and the current date and weather. Other perceptual KSes provide presence and identity information of users in front of the display, through the use of local sensors such as cameras and motion sensors.

2. *Domain-specific expert* Knowledge Sources. The second category of knowledge sources contains domain-specific experts, which contribute external wisdom into the current situation by triggering on knowledge produced by lower-level perceptual KSes as well as by other domain-expert peers. An example of a simple domain-specific expert KS would be an agent which knows all national holidays, and posts these holidays when appropriate days arrive. These agents, in effect, classify concrete states of the world into familiar situational characterizations that are recognizable by the *recommender agents*.

3. *Recommenders*. Recommender agents form the highest-level agents in the k:info blackboard architecture. These agents associate world state with candidate items to display. Once an appropriate combination has been identified, a recommender posts a *recommendation* for either a single specific item, or a broader class of candidate displayable information items. These recommendations include a numeric value, indicating the strength of the recommendation, a reference to the item(s) being recommended, and the name of the recommender agent who made the recommendation. *Case-based* or *collaborative* recommenders use statistical collaborative filtering or classification techniques to make recommendations based upon "learned" past interactions, once sufficient data has been acquired.

3.4 Scheduling Display Items

Scheduling items for display, then, involves collecting the posted recommendations and producing a display schedule. The current simple scheduler bases its schedule on the total recommendation level (calculated as a sum over the recommendations for each item) as the probability that a particular item will be displayed next. Items with negative recommendation totals are omitted from the schedule.

4. FUTURE WORK

4.1 Performance and User Evaluation

The most important work that has yet to be completed is an evaluation of the system. From a developer perspective, the blackboard architecture has provided a useful structure that has simplified system development and improved modularity. A user study is planned, which will survey users as to whether they found displayed items to be of interest, and whether they felt the display provided a useful or detrimental distraction in public spaces.

5. ACKNOWLEDGMENTS

This work is being done with the Agent-based Intelligent Reactive Environments (AIRE) group at the MIT CSAIL Laboratory, under the supervision of Dr. Howard Shrobe and Dr. Kimberle Koile. It is being funded by MIT Project Oxygen.

REFERENCES

[1] D. Cohen and L. Prusak. *In Good Company: How Social Capital Makes Organizations Work*. Harvard Business School Press, Cambridge, Massachusetts, 2001.

[2] J. Laarni. Searching for optimal methods of presenting dynamic text on different types of screens. In *Proceedings of the second Nordic conference on Human-computer interaction*, pages 219–222. ACM Press, 2002.

[3] R. Reddy, L. Erman, R. D. Fennell, and R. B. Neely. The HEARSAY speech understanding system: An example of the recognition process. *IEEE Transactions on Computers*, pages 427–431, 1976.

[4] Employees on the move. *Steelcase Workplace Index Survey*, Apr. 2002.

[5] M. Van Kleek. Intelligent environments for informal public spaces: the Ki/o Kiosk Platform. M.Eng. Thesis, Massachusetts Institute of Technology, Cambridge, MA, February 2003.

[6] T. Winograd. Architectures for context. *Human-Computer Interaction*, 16:401–409, 2001.