Intl. Journal of Human–Computer Interaction, 30: 237–249, 2014

ISSN: 1044-7318 print / 1532-7590 online
DOI: 10.1080/10447318.2013.849547

Multipurpose Public Displays: Can Automated Grouping of Applications and Services Enhance User Experience?

Christos Katsanos[1], Nikolaos Tselios[2], Jorge Goncalves[3], Tomi Juntunen[3], and Vassilis Kostakos[3]
[1]*Hellenic Open University, Patras, Greece*
[2]*University of Patras, Patras, Greece*
[3]*University of Oulu, Oulu, Finland*

Transitioning from bespoke single-purpose displays to multipurpose public interactive displays entails a number of challenges. One challenge is the development of usable mechanisms that allow users to explore the functionality and services on such displays. This article presents a field trial that employs AutoCardSorter, a tool that uses semantic similarity and clustering algorithms, to automatically group the available applications of a public interactive display into categories based on the developer-provided descriptions of each application. The results demonstrate that the grouping generated by AutoCardSorter improved both performance and self-reported usability measures compared to practitioners' existing grouping. In addition, the study investigated the interplay between grouping and interaction modality (i.e., public display vs. desktop). Results tend to support that grouping affects more the user experience with a multipurpose interactive display, but findings were insignificant. This work provides a way for public displays to dynamically update their offered services without sacrificing usability.

1. INTRODUCTION

Research on public interactive displays has come a long way, and recent years have seen a substantial interest on this topic. An open research challenge is managing the transition from single-purpose to multiapplication public interactive displays that may have dozens of applications (Ojala et al., 2012) and associated application distribution systems from which applications and services are dynamically obtained (Clinch, Davies, Kubitza, & Schmidt, 2012). The introduction of these mechanisms has the potential to revolutionize the software and hardware ecosystem of public displays, as was the case with smartphone application stores.

Although it is not clear where one can draw the distinction between single-purpose and multipurpose displays, we argue that one way to make this distinction is to consider if the richness of the display is in terms of information or functionality. Therefore, we can consider a display with multiple types of information about a city as a single-purpose display in the sense that it provides multiple types of information through one single "application" or interface. On the other hand, we consider as multipurpose a display that provides multiple types of functionality, such as information browsing, games, galleries, and polls to name a few. The transition from single-purpose to such multiapplication public interactive displays raises a number of challenges in their design. For instance, it is not clear how the applications should be presented, organized, and browsed by users.

In this article, we argue that by identifying ways in which applications and services on a multipurpose public display can be grouped into categories could substantially improve the user experience. However, generating this grouping manually can be a laborious process, especially if one takes into account the expected update frequency and sheer number of potential applications and services in multipurpose public interactive displays. Thus, we are interested in automating this process to the largest extent possible. To this end, a tool for automatically grouping applications into categories given the developer-provided applications' textual descriptions is presented. In addition, we report a field trial in which we evaluate the output of the tool by measuring its impact on actual and perceived usability through user evaluation. In the context of this field trial, we also investigate the impact of information grouping on user experience with a multipurpose interactive display by comparing usability measures in the desktop and public display modalities.

2. RELATED WORK

Early research on public displays was mostly conducted on single-purpose bespoke public displays, for example, Plasma Posters (Churchill, Nelson, Denoue, & Girgensohn, 2003) or GroupCast (McCarthy, Costa, & Liongosari, 2001). Recent advances in public display technology have enabled increasing numbers of displays to be deployed and installed in public locations. These deployments are increasingly making a transition from static "broadcast" displays to interactive ones.

Address correspondence to Christos Katsanos, School of Science and Technology, Hellenic Open University, 18 Parodos Aristotelous, Patras GR-26335, Greece. E-mail: ckatsanos@eap.gr

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/hihc.

This transition to interactive displays, where members of the public are empowered to control and use the display, has opened a range of new research challenges and at the same time has broadened the design space for public displays (Muller & Kruger, 2006). Whereas on "broadcast" displays the primary challenge is designing for the effective sharing of information with the public, interactive displays' main design requirement is that of interaction: designing and implementing a mechanism for members of the public to browse, navigate, and identify information that the display makes available (Congleton, 2007).

Orthogonal to the challenges introduced by interactivity, the move from single-application to multiapplication public interactive displays opens up a whole new range of problems but also new design space (Ojala et al., 2012). Although most previous work has considered supporting multiple simultaneous users on a single display (Izadi, Brignull, Rodden, Rogers, & Underwood, 2003), it has rather overlooked the issues raised by having users interact with multiple applications on a single public display (Ojala et al., 2012).

Instead, researchers have focused on various *in-app* menu structures in the context of multiuser, multitouch tabletop/surface displays. A number of limitations have been identified (Bailly, Lecolinet, & Guiard, 2010; Brandl et al., 2009; Nacenta, Baudisch, Benko, & Wilson, 2009; Rooney & Ruddle, 2012) including occlusion (the hand and fingers may hide parts of the menu display), accuracy (the large surface area of finger-screen contact may induce item selection error), and reachability (some menu items may be unreachable). Several prototype solutions have been presented to alleviate these concerns, and examples include using gestures (Kostakos & O'Neill, 2003) or multiple modalities (O'Neill, Kaenampornpan, Kostakos, Warr, & Woodgate, 2006); Marking menus (Kurtenbach & Buxton, 1993), which extend the concept of pie menus (Callahan, Hopkins, Weiser, & Shneiderman, 1998) by allowing a user to perform a selection by either popping up a menu or drawing a mark in the direction of the desired item; and Control menus (Pook, Lecolinet, Vaysseix, & Barillot, 2000) that combine command selection and direct manipulation so that users do not have to switch focus between menus and other interactors.

Beyond these usability concerns at the *lexical* level of interaction, multipurpose public displays raise usability concerns at the *syntactic* level (Dix, Finlay, Abowd, & Beale, 2003): Users need to browse through the various applications or services of the display before activating any one of them. This is effectively a task of searching information in the context of a public interactive display and is likely the first task of every user of such a public display. Therefore, it is important that the top-level application grouping of the display is efficient and does not cause frustration or fatigue.

However, the design of user-centered information architectures is challenging (Jones et al., 2006; Toms, 2002). A variety of methods, such as card-sorting, contextual inquiry and ethnographic interviews, are available, but the transition from descriptive user research models to prescriptive design remains challenging (Sinha & Boutelle, 2004). Given the expected update frequency and sheer number of potential applications and services in public multipurpose displays, an effective automated approach can greatly contribute to the development of usable top-level application groupings.

In the context of desktop computing, improper information architecture has been shown to cause various usability problems and deteriorate the overall interaction experience. For instance, users may have difficulty forming a cognitive model of the information structure (Otter & Johnson, 2000) and can become lost because of the nonlinear nature of hypermedia navigation (Chen & Macredie, 2002). Despite the increasing popularity of public displays, no work has considered how information structure and grouping of applications and services affect usability on public displays. We attribute this to the initial popularity of bespoke public displays, where application grouping was simply not a concern. However, we cannot assume that information structures that work on desktops will be equally effective in other interaction contexts (Kim, Jacko, & Salvendy, 2011), such as multipurpose public displays. Currently, the impact of application grouping on findability (Morville, 2005) remains unknown in the context of multipurpose public display interaction, and at best only assumptions can inform design.

In summary, public displays are increasingly becoming multipurpose, and the development of application distribution systems for public displays suggests that the diversity and number of applications on these systems is likely to grow. This raises the challenge of helping users browse through the available applications and launch the one they are interested in. However, due to the lack of automated grouping mechanisms, so far applications had to be grouped manually, thus leading to increased human effort and possible inefficiencies (Katsanos, Tselios, & Avouris, 2008a, 2008b). In addition, the impact of application grouping on user's experience with a multipurpose display is rather unexplored compared to the rich available literature for the web (Chen & Macredie, 2002; Otter & Johnson, 2000; Patsula, Detenber, & Theng, 2010; Sinha & Boutelle, 2004), desktop computing (Jacko, Salvendy, & Koubek, 1995; Kim et al., 2011; D. P. Miller, 1981; Smelcer & Walker, 1993) and mobile phones (Chae & Kim, 2004; Geven, Sefelin, & Tscheligi, 2006; Kim et al., 2011; Ling, Hwang, & Salvendy, 2007).

3. AN AUTOMATED TOOL FOR APPLICATION GROUPING

In prior work (Katsanos et al., 2008a, 2008b), we have proposed AutoCardSorter, a tool that automates grouping of web pages in the context of website structuring, and demonstrated that it simulates well actual card-sorting sessions. Our findings were based on three independent studies that compared participants' card-sorting groupings of web pages (treated as documents) against the ones produced by the tool. In all studies we reported that the groupings provided by the tool significantly correlate with those of card-sorting participants ($r = .52$–$.80$, $p < .01$) and result in highly similar website organization schemes (Katsanos et al., 2008b). Further reporting of data on the tool's ability to simulate card-sorting sessions is beyond the scope of this article.

However, so far we have no validation of the tool's performance when working with extremely short-text documents, as are the application descriptions on a multipurpose public display. Typically these are one to three sentences, and thus they offer a rather restrictive context within which our tool needs to perform. Furthermore, we have no empirical evidence on the effect that the tool's output has on usability, and whether even small discrepancies in the groupings sacrifice usability. In this article, we extend previous work by describing a study that investigates the tool's performance in a novel interaction context (i.e., multipurpose public displays) that poses significant challenges, such as the very brief textual descriptions of the applications, and considers the tool output's effect on actual and perceived usability.

In this section, we briefly describe the AutoCardSorter grouping algorithm (for an elaboration, see Katsanos et al., 2008b) and introduce a new version of the tool in the context of application grouping on a multipurpose public display.

3.1. AutoCardSorter Algorithm and Typical Usage Scenario

The tool takes as input a list of applications and their respective textual descriptions and generates a proposed grouping for these applications. To achieve this, AutoCardSorter relies on semantic similarity measures, in specific Latent Semantic Analysis (LSA; Landauer, McNamara, Dennis, & Kintsch, 2007), and clustering algorithms (Jain, Murty, & Flynn, 1999) to simulate an open card-sorting exercise. Our tool is addressed at practitioners responsible for the design, maintenance, and evaluation of multipurpose public displays. Currently, AutoCardSorter is made freely available upon request.

Treating each application description as a "document," the tool uses the document-to-document version of LSA to produce a matrix of semantic similarities between each pair of the developer-provided applications' textual descriptions. LSA is "a mathematical method for computer modeling and simulation of the meaning of words and passages by analysis of representative corpora of natural text" (Landauer & Dumais, 2008, para. 1). The output of LSA is a value measuring the semantic similarity between a given pair of documents. Each LSA value lies between +1 (identical) and –1 (opposite), with near-zero values representing unrelated documents. An LSA value for a pair of documents is meaningful only for a specific semantic space, that is, for a large body of text used to train the algorithm. This training text is assumed to represent the reading and understanding skills of the population of interest, for example, general reading up to 1st-year college.

Given the matrix of semantic similarities between all possible pairs of application descriptions, AutoCardSorter subsequently employs a hierarchical agglomerative clustering algorithm to produce groupings of applications based on the semantic similarity of their descriptions. Clustering algorithms address the problem of assigning a set of objects into groups so that objects that are placed in the same group are more similar according to a criterion (Jain et al., 1999; Witten & Frank, 1999). In our context, applications that are placed in the same group share a higher degree of semantic similarity (i.e., LSA value), compared to applications placed in other groups. The tool offers three agglomerative clustering algorithms (a) average linkage, (b) complete linkage or farthest neighbor, and (c) single linkage or nearest neighbor. Complete linkage tends to produce tightly bound or compact groupings, whereas single linkage has a tendency to produce groupings that are straggly or elongated (Jain et al., 1999). The average linkage is a compromise between the two and usually produces balanced groupings that are easier to interpret (Witten & Frank, 1999).

A typical usage scenario of our tool is the following. First, a set of textual descriptions is provided for all the applications that the public display offers (see Figure 1). To this end, the full text or a selected subset of sentences in the developer-provided descriptions can be used. Then, the practitioner selects an appropriate LSA semantic space to represent the typical end-users of the public display and indicates the desired type of clustering algorithm. Multipurpose public displays are typically addressed to anyone, thus a general reading LSA semantic space should be suitable in almost all cases (default option in tool). In addition, our experience shows that the average-linkage clustering algorithm (default option in tool) typically produces more appropriate groupings compared to the single- and complete-linkage variations.

Then the tool is instructed to run an automated analysis, combining the LSA method with the selected clustering algorithm to generate groupings of applications that are semantically relevant. The output of this analysis is an interactive dendrogram (see Figure 1), which reflects the similarities between applications and supports the creation of potential categories. The tool can be instructed to automatically determine the optimal number of categories by attempting to maximize the variance explained by any particular grouping. However, it is also possible to indicate a desired number of categories either by dragging the line depicting the similarity strength among the grouped items or by explicitly specifying the desired number of categories and the tool will generate as many.

4. STUDY

The goal of the study is to establish whether the application grouping generated automatically with AutoCardSorter (Tool Grouping) can improve the navigation structure and thus the user experience over an application grouping generated by practitioners (Practitioner Grouping). In addition, the study

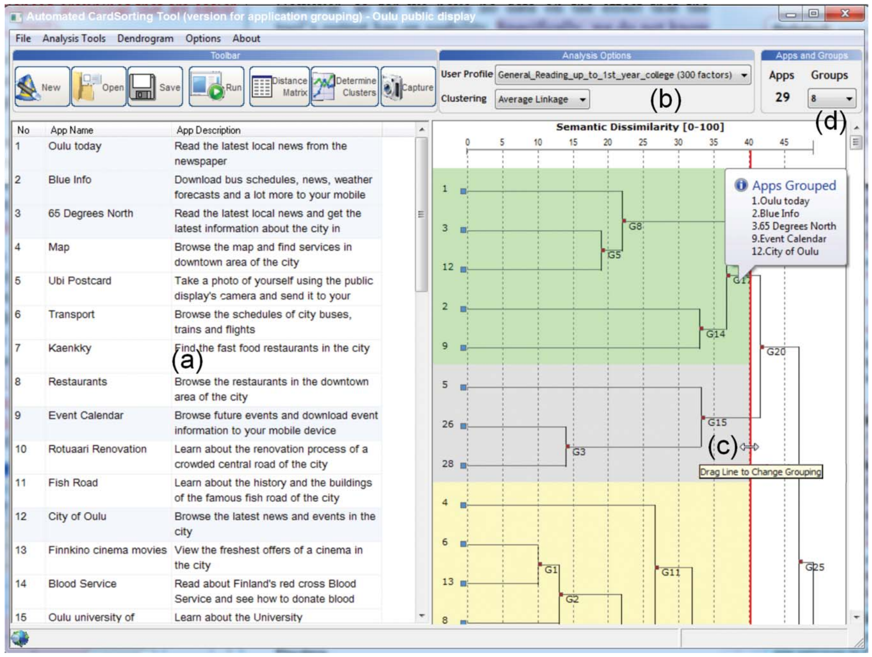

FIG. 1. Screenshot of AutoCardSorter (version for application grouping) and its main functionalities. (a) Applications' textual descriptions, (b) analysis options with default values, (c) interactive dendrogram depicting potential groupings, (d) specifying desired number of top-level groupings.

aims to investigate the impact of information grouping on user experience with a multipurpose interactive display by comparing usability measures in the desktop and public display modalities. We use the term "public display modality" to refer to touch-based interaction with a public large screen display, whereas the term "desktop modality" implies mouse-driven interaction with a personal much smaller display.

We conducted a field study in the context of an ongoing longitudinal deployment of a multipurpose public interactive display deployed in a busy lobby of a university. This display provides 29 applications that had been grouped into eight categories by practitioners. Using our tool a new grouping of the applications was derived while keeping constant the number of categories. We hypothesize that when compared to the practitioners' application grouping, the automated grouping will improve usability, thus leading to

- increased *effectiveness*, as measured by task completion rate;
- increased *efficiency*, as measured by task completion time, number of clicks, and the lostness metric (Smith, 1996). In the context of hypertext navigation, the lostness metric was defined (Smith, 1996) as:

$$L = \mathrm{sqrt}\left[\left(\frac{N}{S} - 1\right)^2 + \left(\frac{R}{N} - 1\right)^2\right]$$

  where R is the minimum number of page visits in an 'ideal' situation, N is the number of different pages visited and S is the total number of page visits including revisits. A perfect search has $L = 0$ and L increases as lostness increases. Smith (1996) suggested that when the L value exceeds 0.5 the user is definitely lost and a closer examination is required when the metric lies between 0.4 and 0.5;
- increased *subjective assessments of usability*, as measured by a posttask, 7-point question for confidence in item choice averaged across all given tasks, and two postexperiment scales: the 10-item System Usability Scale (SUS; Brooke, 1996) and the single-item, seven-point adjective rating scale (Bangor, Kortun, & Miller, 2008, 2009).

4.1. Method

Design

The study was a 2 × 2 experimental design. The first independent variable was the application grouping (one created by practitioners, another by our tool), and the second independent variable was interaction modality (touch-based interaction with a public large screen display, mouse-driven interaction with a personal much smaller display). Participants were randomly allocated to two nonoverlapping conditions, counterbalancing the order of the conditions. This allowed us to complete more trials per participant without introducing bias in our results. All trials took place in the same public setting to address any location-based bias.

Participants

In total 40 participants (12 female, 28 male) aged 18 to 41 ($M = 27.6$, $SD = 5.6$) participated in the study. They were recruited randomly in situ in front of a multipurpose public interactive display deployed in a busy lobby of a university. All participants were informed about the purpose and duration of the experiment and were rewarded with a voucher for the local cafeteria. Some participants had prior exposure to the public display and practitioner grouping conditions, but as we report in the results this did not have an effect.

Apparatus

A 57-in. touch screen was used as the public display, mounted vertically at 1.2 meters from the floor. A 12.5-in. Lenovo X220 laptop with a three-key external mouse was used as the PC, placed on a desk adjacent to the public display. One grouping was produced by a team of three practitioners with accumulated experience in software engineering and UX methods. The practitioners team was responsible for the design and maintenance of the public display and followed card-sorting and field trials to produce an initial application grouping, which was later modified as new apps became available during its 2-year deployment. This reflects a common real-world case of a grouping that emerged over time and constituted the practitioner grouping condition in our study. A software engineer experienced in interaction design and unfamiliar with both the applications and the practitioner grouping used our tool to generate an alternative grouping (tool grouping). This was achieved using the "General Reading up to 1st-year college" semantic space, and the average linkage clustering algorithm. The category names were derived manually. Both Groupings had eight categories and the same 29 applications. Table 1 presents the two groupings.

Participants interacted with a portal (see Figure 2) that presented the application categories and their associated applications. The portal adopted either the practitioner or tool grouping, and the same portal was used on both the public display and desktop. Participants could click on application categories to display the associated applications and their textual descriptions (see Figure 2). Participants could then click on an application itself and press a button to indicate that they wished to use it.

Procedure

Each participant initially received training on the system. A printed screenshot of the portal was given to them, and they received training on the use of application categories and their associated applications. Participants were instructed to find the correct application for each of seven tasks. To choose an application they simply had to click on it and confirm their choice by pressing an "OK" button. Examples of tasks given to participants are the following (in parentheses the developer-provided

TABLE 1
The Grouping as Derived by the Practitioners (Left) and the Automated Tool (Right)

Practitioner Grouping	Tool Grouping
News	*News & Events*
1.1 Oulu today	1.1 Oulu today
1.2 Blue Info	1.3 65 Degrees North
1.3 65 Degrees North	3.5 City of Oulu
	3.2 Event Calendar
	1.2 Blue Info
Services	*Games & Fun*
2.1 Map	5.1 UBI Mosquitos
2.2 UBI Postcard	5.2 Ubitris
2.3 Transport	5.3 Belle Memory
2.4 Kaenkky	5.4 Hangman
	8.4 Fun Square
City	*City Guide*
3.1 Restaurants	2.1 Map
3.2 Event Calendar	2.3 Transport
3.3 Rotuaari Renovation	4.1 Finnkino cinema movies
3.4 Fish Road	2.4 Kaenkky
3.5 City of Oulu	3.1 Restaurants
3rd party	*3rd party services*
4.1 Finnkino cinema movies	4.2 Blood Service
4.2 Blood Service	4.3 Oulu university
4.3 Oulu university	4.4 Etuovi.com
4.4 Etuovi.com	8.2 Digifields
4.5 Media sales	4.5 Media sales
Fun and games	*Socialize & Share*
5.1 UBI Mosquitos	8.1 Run With Us
5.2 Ubitris	8.3 CLIO
5.3 Belle Memory	2.2 UBI Postcard
5.4 Hangman	
Multimedia	*Feedback & Surveys*
6.1 Kihina meets UBI	7.1 Survey
6.2 Streetgallery	7.2 UBI questionnaire
Help and survey	*Oulu History*
7.1 Survey	3.3 Rotuaari Renovation
7.2 UBI questionnaire	3.4 Fish Road
New cool stuff	*Art & Culture*
8.1 Run With Us	6.1 Kihina meets UBI
8.2 Digifields	6.2 Streetgallery
8.3 CLIO	
8.4 Fun Square	

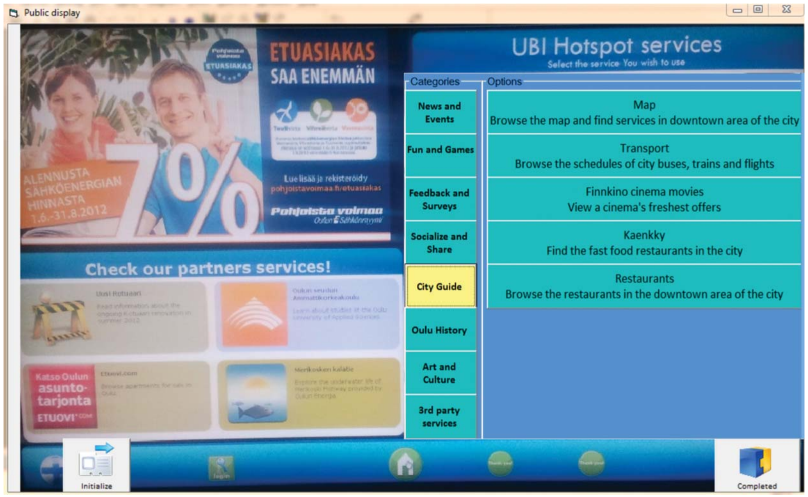

FIG. 2. Screenshot of the portal that presented the application categories and their associated applications to participants.

FIG. 3. The setting of the field trial. Left: A participant using the public display and a passerby. Right: A participant using the PC.

name of the target-application): "*Find the bus schedules in the city of Oulu*" (Transport), "*You know that Oulu is famous for the mosquitos? Play a game trying to kill some of them*" (UBI Mosquitos), and "*Find an application that will allow you to post a photograph on Facebook*" (UBI Postcard). The task order was randomized. Figure 3 presents the environment in which the experiment took place.

The portal logged the following data during trials: *participant id, condition, trial number, selected category, selected application,* and *time*. After each completed task, participants were asked to rate their confidence about selecting the correct application on a 7-point discrete visual analog scale with one question: "How confident do you feel that you selected the right item?" Upon completing all tasks they answered a questionnaire with demographics (age, years of Internet use, prior exposure to the public display), a standardized adjective satisfaction scale with one question (Bangor et al., 2008, 2009), and the SUS with 10 questions (Brooke, 1996). Finally, we performed semistructured interviews with the participants in order to get their views on both used groupings. Each participant took about 20 min to complete all tasks and questionnaires.

4.2. Results

Table 2 contains statistics for the dependent variables grouped per condition. We also provide aggregate statistics per grouping and modality separately. Overall, the tool grouping improved both participants' performance and perception. In addition, participants using the public display were more hesitant (less clicks) and tried to finish faster but reported to like the public device more.

Prior exposure and usage frequency of the public display were investigated as potential covariates. Analyses showed no significant effects for both variables, thus no adjustments in the means of the dependent variables were required. Mixed-effect models were then conducted to investigate for any main or interaction effects of grouping and modality on the collected measures. Effect sizes were calculated according to Field (2009).

Performance Measures

Task performance measures include task success rate, task completion time, number of clicks, and the lostness (Smith, 1996) metric. Analyses were conducted using composite scores

TABLE 2
Participants' Task Success, Task Completion Time, Number of Clicks, Lostness Score, Confidence in Choice, Adjective Rating and SUS Score in Relation to the Application Grouping and Interaction Modality

		Task Success (%)		Task Completion Time (s)		No. of Clicks		Lostness Score		Confidence in Choice (1–7)		Adjective Rating (1–7)		SUS Score (0–100)	
	N	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD	M	SD
Practitioner Grouping & Public display	20	81.2	18.0	18.6	8.3	4.0	1.3	0.35	0.14	5.9	1.1	4.5	1.1	67.4	13.3
Tool Grouping & Public Display	20	92.8	9.8	12.1	6.2	3.3	0.9	0.24	0.12	6.2	1.0	5.1	0.7	75.0	13.5
Tool Grouping & Desktop	20	91.6	12.3	15.4	7.0	3.5	1.1	0.29	0.16	6.2	1.4	4.8	1.2	72.1	18.0
Practitioner Grouping & Desktop	20	87.8	15.6	17.2	9.3	5.3	1.5	0.44	0.13	6.0	0.9	4.1	1.2	58.5	22.0
Practitioner Grouping	40	84.5*	17.0	17.9*	8.7	4.6***	1.5	0.39***	0.14	5.9	1.0	4.3**	1.1	62.9**	18.5
Tool Grouping	40	92.2*	11.0	13.7*	6.7	3.4***	1.0	0.27***	0.15	6.2	1.2	4.9**	1.0	73.6**	15.7
Desktop	40	89.7	14.0	16.3	8.8	4.4*	1.6	0.36*	0.17	6.1	1.2	4.4	1.2	65.3	21.0
Public Display	40	87.0	15.5	15.3	7.9	3.7*	1.2	0.30*	0.14	6.0	1.1	4.8	0.9	71.2	13.8

Note: Task success, task completion time, number of clicks, and confidence in choice are scores averaged across all the seven tasks used in the study.
$^{*}p < .05$. $^{**}p < .01$. $^{***}p < .001$.

over all seven application-finding tasks for all measures: task success as the percentage of tasks in which the participant found the correct item, and task completion time, number of clicks, and lostness score as the mean value of each corresponding measure for all tasks.

Task success. The effect of grouping on task success was significant, $F(1, 74) = 6.73$, $p = .011$, $r = .29$. Participants were significantly more successful when using the tool grouping. By contrast, no significant effect of modality on task success was found, $F(1, 74) = 0.83$, $p = .366$, $r = .11$. There was also no interaction effect of grouping and modality on task success, $F(1, 74) = 2.48$, $p = .120$, $r = .18$.

Task completion time. Grouping had a significant effect on task completion time, $F(1, 74) = 6.87$, $p = .011$, $r = .29$. The tool grouping led to significantly faster times for task completion. There was no significant effect of modality on task completion time, $F(1, 74) = 0.36$, $p = .548$, $r = .07$. The interaction of grouping and modality with respect to task completion time was not significant, $F(1, 74) = 3.07$, $p = .084$, $r = .20$.

We note that participants using the public display improved their task completion times more when they moved from the practitioner grouping to the tool grouping (see Figure 4). Given that all order effects were counterbalanced, this finding could suggest that in public displays the quality of the information structure is more important for achieving task efficiency. However, the latter finding should be handled with care, as it was insignificant ($p = .084$) and had a small (Cohen, 1992) effect size.

Number of clicks. The effect of grouping on participants' number of clicks was significant, $F(1, 74) = 20.83$, $p < .001$, $r = .47$. The tool grouping led to significantly fewer clicks.

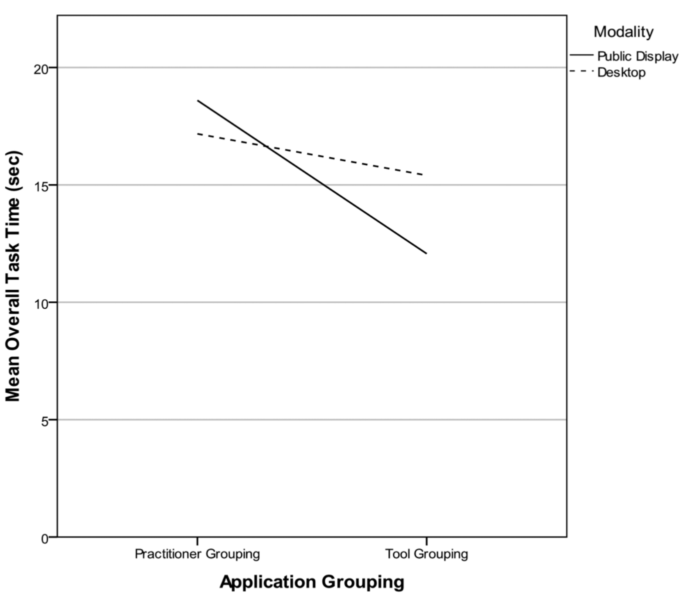

FIG. 4. The effect of application grouping and modality on participants' task completion time.

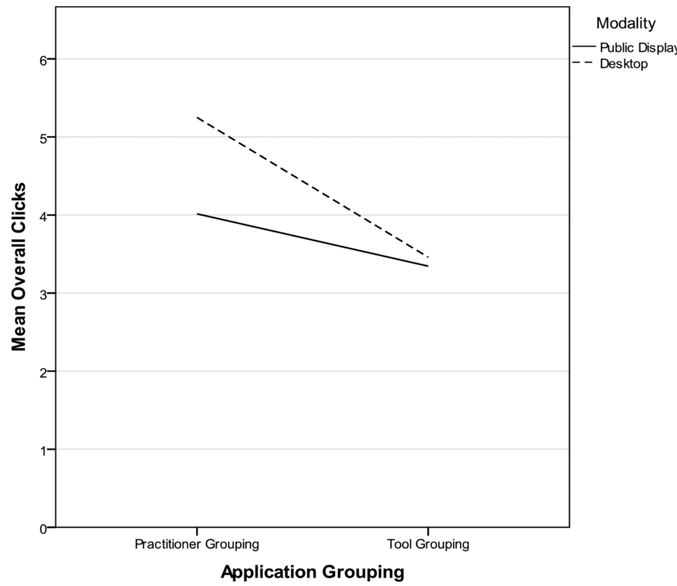

FIG. 5. The effect of application grouping and modality on participants' number of clicks.

Modality had also a significant effect on participants' number of clicks, $F(1, 74) = 6.27$, $p = .014$, $r = .28$. Participants using the public display made significantly fewer clicks than when they used the desktop. The interaction effect of grouping and modality on participants' number of clicks was borderline-insignificant, $F(1, 74) = 3.96$, $p = .050$, $r = .23$.

As Figure 5 shows, the decrease in the number of clicks was steeper for the desktop modality. Participants using the public display might have followed an optimizing selection strategy, instead of a satisficing one, either because of feeling observed by passersby or because they wanted to minimize the movement associated with their selections. However, this interaction effect requires additional investigation because in this study it was borderline-insignificant ($p = .050$) and had a small to medium (Cohen, 1992) effect size.

Lostness score. Grouping had a significant effect on the lostness score, which captures the pattern of item visits during tasks, $F(1, 74) = 15.15$, $p < .001$, $r = .41$. The tool grouping led to significantly less disoriented participants. A significant effect of modality on the lostness score was also found, $F(1, 74) = 4.68$, $p = .034$, $r = .24$. Participants using the public display appeared to be significantly less disoriented than when they used the desktop. The interaction of grouping and modality with respect to participants' lostness score was not significant, $F(1, 74) = 0.34$, $p = .563$, $r = .07$.

Self-Reported Measures

Self-reported measures were collected using a posttask question for confidence in item choice, and two postexperiment scales: the adjective rating scale proposed by Bangor et al. (2008, 2009), and the SUS (Brooke, 1996). Posttask user ratings for confidence in item choice were averaged across all tasks to produce a composite score that was used in the analyses.

Posttask confidence rating. The effect of grouping on posttask confidence rating was not significant, $F(1, 74) = 1.45$, $p = .233$, $r = .14$. The effect of modality was also not significant, $F(1, 74) = 0.06$, $p = .807$, $r = .03$. No interaction effect of grouping and modality on posttask confidence in participants' choices was observed, $F(1, 74) = 0.34$, $p = .563$, $r = .07$.

Postexperiment adjective rating. Grouping had a significant effect on adjective ratings, $F(1, 74) = 7.69$, $p = .007$, $r = .31$. Participants gave significantly higher ratings when using the tool grouping. Analysis did not show a significant main effect of modality, $F(1, 74) = 2.23$, $p = .140$, $r = .17$, or a significant interaction effect, $F(1, 74) = 0.03$, $p = .866$, $r = .02$.

Postexperiment SUS score. The effect of grouping on SUS scores was significant, $F(1, 74) = 7.64$, $p = .007$, $r = .31$. When using the tool grouping, participants gave SUS scores that were on average 10 points higher. By contrast, no significant effect of modality on SUS score was found, $F(1, 74) = 2.34$, $p = .131$, $r = .18$. There was also no interaction effect of grouping and modality, $F(1, 74) = 0.525$, $p = .471$, $r = .08$.

5. DISCUSSION

Recent advances in technology have resulted in the deployment of an increasing number of displays in public locations. These deployments have successfully made a transition from static "broadcast" displays to interactive ones. Until today, usability research on public interactive displays has tended to focus on interaction at the lexical level, for instance, considering hand occlusion (Bailly et al., 2010), click accuracy (Brandl et al., 2009), and reachability of screen items (Nacenta et al., 2009). However, the move from single-application to multiapplication public interactive displays opens up a whole new range of problems at the syntactic level, as users now have to figure out the functionality and applications available on the display and choose one of them (Hosio, Goncalves, & Kostakos, 2013; Kostakos, Kukka, Goncalves, Tselios, & Ojala, 2013).

Of interest, previous work has considered supporting multiple simultaneous users on a single display but completely overlooked the issues raised by having users interact with multiple applications on a single public display. We argue that the transition to multipurpose displays raises a cascade of new questions. For example, how should the multiple applications be presented to users? How many applications should a display have? Should users be able to install their own applications on the displays? All these are questions that we hope to answer in our future work. For now, however, our research has some important implications in terms of application grouping in multipurpose public displays.

The main finding is that the proposed automated categorization technique produced groupings in an efficient manner without sacrificing actual and perceived usability. In specific, all four task performance measures and two out of three of the self-reported measures were significantly higher in the grouping of applications produced by the tool compared to the existing grouping produced by practitioners. We also investigated the impact of grouping on users' performance and perceptions while interacting with a multipurpose public display, which is rather unexplored compared to the literature (see Kim et al., 2011, for a review) in the desktop or mobile interaction paradigm. Results tend to support that grouping affects more the user experience with a multipurpose interactive display, but findings were insignificant. Based on both quantitative and qualitative data, we tentatively argue that users appear to change browsing behavior, which could explain the aforementioned tendency.

5.1. Automating Application Grouping Without Sacrificing Usability

Our study shows that all task performance measures (success, time, number of clicks, lostness) were significantly improved by the tool's grouping of applications. Moreover, the participants' perceived usability assessments were significantly higher for both the adjective rating and the SUS score. This is an important finding, as to the best of our knowledge no other study examined the impact of different groupings on perceived usability. Because users' perceptions greatly influence adoption and usage of computational technology, this finding reinforces the importance of designing a proper information structure.

On the contrary, the participants' posttask confidence in choice was not significantly affected by application grouping. The collected data showed that in most cases participants managed to find the correct application in both groupings, but at the expense of time and clicks while interacting with the practitioner grouping. This appears to have significantly influenced their perceived usability assessments but not their confidence in their selection.

It is worth noting that the majority of the participants (29 of 40) had prior exposure to the multipurpose public display, which presented the applications using the practitioner grouping. Despite this "handicap," the tool grouping outperformed the practitioner grouping, even in the short duration of our field trial. Therefore, we can expect that over longer periods the tool grouping could result in even greater usability gains. Also, the reported differences may become amplified if a three-tiered menu structure is adopted, as incorrect selections at the top and the second level impose additional click costs compared to two-tiered structures (C. S. Miller & Remington, 2004). Moreover, structures with up to eight items per tier produce faster results (C. S. Miller & Remington, 2004); therefore, the effect of grouping on bigger structures is expected to be more substantial in comparison to the results reported in this article.

The majority of participants reported in the interviews that the tool grouping was more intuitive, with categories being more consistent and logical making it easier to find the correct application. The participants were not informed about which navigation structure was the original and which structure the tool created. During the interviews some participants highlighted a (false) sense of familiarity with the tool grouping, claiming that they already had experience with it. This was clearly not the case, but nevertheless the tool grouping felt more familiar to some participants.

5.2. Grouping Tends to Affect More Interaction on Public Displays Than on Desktop Computing

When interpreting the results by taking into account the different modalities (i.e., public display vs. desktop), the striking difference lies in the significantly fewer number of clicks on the public display setting. Moreover, participants using the desktop decreased more the number of clicks when they moved from the practitioner grouping to the tool grouping (Figure 5). This is also reflected by the lostness metric, which in the practitioner grouping was at alarming (Smith, 1996) rates (0.44) on the desktop while significantly lower on the public display. Moreover, participants using the public display improved their task completion times more when they moved from the initial application grouping to the tool-based one (Figure 4). These differences, may suggest a change in participants' browsing behavior while using the public display.

First, because their behavior could be observed by passersby, some participants might have shown a more "careful" behavior by examining more thoroughly the available options. In fact, some participants reported in the follow-up interviews that they became frustrated when they were not sure that they performed the task correctly on the public display while the same did not happen on the desktop. The latter is in alignment with existing literature (Kaviani, Finke, Fels, Lea, & Wang, 2009; Vogel & Balakrishnan, 2004) reporting on the privacy concerns of users while interacting with public displays. In addition, because they are standing while interacting with the public display, they possibly tended to minimize their movement, which in turn is reflected on significantly lower number of actions. Furthermore, they could have felt more comfortable in the desktop setting, thus they perceived the cost of extra clicks as negligible. The differences in strategies adopted by the participants were very evident during most trials. In general, participants, while using the desktop, seemed less concerned with completing the task successfully on their first tries and browsed in quick succession across multiple categories.

When other nondesktop contexts are concerned, Geven et al. (2006) demonstrated that narrow hierarchies perform better than the broader hierarchies on mobile devices, contrary to studies for desktop use (Larson & Czerwinski, 1998; D. P. Miller, 1981) that suggest keeping the number of levels in the information hierarchy to a minimum. Nevertheless, a broad hierarchy will be always beneficial for frequent and experienced users, as by definition it requires a reduced number of clicks. However, such an expansion in breadth might have its limit, as C. S. Miller and Remington (2004) illustrated. In our study, no category contained more than five items in both groupings; therefore an effect of this parameter on the obtained results is expected to be negligible. However, the findings provided by Geven et al. reinforce the concern whether there are differences in users' behavior away from the desktop. As a result, further studies on information grouping for public displays are required to shed light and provide conclusive evidence.

5.3. Challenges and Limitations

Müller argued that public displays do not invite people for a single reason, but users come across them with no dedicated purpose (Müller, Alt, Michelis, & Schmidt, 2010). Also, when a display features multiple applications, many applications are launched due to curiosity or play rather than intention of using the application (Hosio, Kukka, Jurmu, Ojala, & Riekki, 2010). These issues have consequences when designing and evaluating multipurpose displays.

We argue that the design of a multipurpose public display should focus on long-term desirability. Integration of various features and applications for such a diverse group of users constitutes a challenging task that is further complicated if subsequent feature additions are taken into account. Hence an important challenge will be to manage the constant update of the functionality on such public displays.

Our work suggests that an automated tool can be useful in identifying ways to group an ever-increasing set of applications, but we believe that we are far from fully automating this process. First, in our study the application category names were generated by the user of the tool, not the tool itself. It is not clear whether this task can be easily automated. This category-naming process may have affected our results, but all application names and application descriptions were identical across all conditions. Because desktop-related literature suggests that the impact of high-quality labels on users' performance is significant (Katsanos, Tselios, & Avouris, 2010a; Resnick & Sanchez, 2004; Tselios, Katsanos, & Avouris, 2009), more studies are required to investigate the extent of significance to multipurpose public interactive displays.

Another challenge in automating grouping construction is managing the number of groups. In our case the tool was used to generate eight categories simply because this was the number of categories already present in the actual system. However, a characteristic of automated clustering is that the addition of a small number of new applications may lead to very different clusters and groups of applications. Nevertheless, as Danielson (2003) suggested, it is advised to keep the same organization structure while adding more applications. Applications that change grouping in subsequent system's revisions are likely to impact their findability and increase user frustration and disorientation. However, a new grouping of applications might be necessary if a certain number of applications do not seem to belong to one of the existing groups or, to put it in other words, have weak information scent (Jones et al., 2006).

Hence, we do not argue that the tool we have presented can fully replace a human administrator. Rather, we believe that the tool should be integrated in the design process, similarly to automated tools that have been proposed in the context of website design (see Katsanos, Tselios, & Avouris, 2010b, for a review), and should be used for generating "advice" to administrators. This is because there are a number of considerations that the tool currently does not take into account, such as the previous grouping, the adverse effects of frequently changing the groups, and the adverse effects of radically changing the groups.

6. CONCLUSION

The findings of this study give valuable insight given the lack of related studies on how information structure and grouping of applications and services affect usability of multipurpose public interactive displays. It was demonstrated that the tool-based grouping increases participants' effectiveness, efficiency, and perceived usability while interacting with a public display, which in turn can lead to increased adoption and frequency of usage. At the same time the proposed automated approach is straightforward and requires minimum resources, and thus can result in its wider adoption, which in turn means improved interaction experience for more users of multipurpose public displays. Toward this end, our efforts are focused on two important additional features of the tool, which are not implemented currently. First, the automatic generation of category names based on keyword selection and, second, the real-time classification of applications as they are introduced to the display, coupled with sufficient feedback to inform frequent users about their existence. Finally, this article provided initial support that information grouping is even more important in the context of multipurpose interactive displays compared to desktop computing. However, further studies are required to provide conclusive evidence on the latter.

REFERENCES

Bailly, G., Lecolinet, E., & Guiard, Y. (2010). Finger-count & radial-stroke shortcuts: two techniques for augmenting linear menus on multi-touch surfaces. In *CHI 2010: Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 591–594). New York, NY: ACM Press.

Bangor, A., Kortum, P., & Miller, J. (2008). An empirical evaluation of the system usability scale. *International Journal of Human–Computer Interaction*, *24*, 574–594.

Bangor, A., Kortun, P., & Miller, J. (2009). Determining what individual SUS scores mean: Adding an adjective rating scale. *Journal of Usability Studies*, *4*, 114–123.

Brandl, P., Haller, M., Leitner, J., Doray, B., Seifried, T., & To, P. (2009). Occlusion-aware menu design for digital tabletops. In *CHI 2009: Ext. Abstracts of the SIGCHI conference on Human Factors in Computing Systems* (pp. 3223–3228). New York, NY: ACM Press.

Brooke, J. (1996). SUS: A quick and dirty usability scale. In P. W. Jordan, B. Thomas, B. A. Weerdmeester, & I. L. McClelland (Eds.), *Usability evaluation in industry*. London (pp. 189–194). UK: Taylor & Francis.

Callahan, J., Hopkins, D., Weiser, M., & Shneiderman, B. (1998). An empirical comparison of pie vs. linear menus. In *CHI 1998: Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 95–100). New York, NY: ACM Press.

Chae, M., & Kim, J. (2004). Do size and structure matter to mobile users? An empirical study of the effects of screen size, information structure, and task complexity on user activities with standard web phones. *Behaviour and Information Technology*, *23*, 165–181.

Chen, S. Y., & Macredie, R. D. (2002). Cognitive style and hypermedia navigation: Development of a learning model. *Journal of the American Society for Information Science and Technology*, *53*, 3–15.

Churchill, E. F., Nelson, L., Denoue, L., & Girgensohn, A. (2003). The plasma poster network: Posting multimedia content in public places. In *INTERACT 2003: Proceedings of the 9th IFIP TC13 Conference on Human–Computer Interaction* (pp. 599–606). New York, NY: IOS Press.

Clinch, S., Davies, N., Kubitza, T., & Schmidt, A. (2012). Designing application stores for public display networks. In *PerDis 2012: Proceedings of the International Symposium on Pervasive Displays* (Article 10). New York, NY: ACM Press.

Cohen, J. (1992). A power primer. *Psychological Bulletin*, *112*, 155–159.

Congleton, B. (2007). Prospero—A "visual commons" framework for public displays. [Summary project report].

Danielson, D. R. (2003). Transitional volatility in web navigation. *IT & Society*, *1*, 131–158.

Dix, A., Finlay, J., Abowd, G. D., & Beale, R. (2003). *Human–computer interaction*. Englewood Cliffs, NJ: Prentice Hall.

Field, A. P. (2009). *Discovering statistics using SPSS*. Thousand Oaks, CA: Sage.

Geven, A., Sefelin, R., & Tscheligi, M. (2006). Depth and breadth away from the desktop: the optimal information hierarchy for mobile use. In *MobileHCI 2006: Proceedings of the 8th International Conference on Human Computer Interaction with Mobile Devices and Services* (pp. 157–164). New York, NY: ACM Press.

Hosio, S., Goncalves, J., & Kostakos, V. (2013). Application discoverability on multipurpose public displays: popularity comes at a price. In *PerDis 2013: Proceedings of the 2nd ACM International Symposium on Pervasive Displays* (pp. 31–36). New York, NY: ACM Press.

Hosio, S., Kukka, H., Jurmu, M., Ojala, T., & Riekki, J. (2010). Enhancing interactive public displays with social networking services. *MUM 2010: Proceedings of the 9th International Conference on Mobile and Ubiquitous Multimedia* (Article 23).

Izadi, S., Brignull, H., Rodden, T., Rogers, Y., & Underwood, M. (2003). Dynamo: A public interactive surface supporting the cooperative sharing and exchange of media. In *UIST 2003: Proceedings of the 26th ACM Symposium on User Interface Software and Technology* (pp. 159–168). New York, NY: ACM Press.

Jacko, J. A., Salvendy, G., & Koubek, R. J. (1995). Modelling of menu design in computerized work. *Interacting with Computers*, *7*, 304–330.

Jain, A. K., Murty, M. N., & Flynn, P. J. (1999). Data clustering: A review. *ACM Computing Surveys*, *31*, 264–323.

Jones, W., Pirolli, P., Card, S. K., Fidel, R., Gershon, N., Morville, P., . . . Russell, D. M. (2006). "It's about the information stupid!": Why we need a separate field of human–information interaction. In *CHI 2006: Ext. Abstracts of the SIGCHI conference on Human Factors in Computing Systems* (pp. 65–68). New York, NY: ACM Press.

Katsanos, C., Tselios, N., & Avouris, N. (2008a). AutoCardSorter: designing the information architecture of a web site using latent semantic analysis. In *CHI 2008: Proceedings of the SIGCHI conference on Human Factors in Computing Systems* (pp. 875–878). New York, NY: ACM Press.

Katsanos, C., Tselios, N., & Avouris, N. (2008b). Automated semantic elaboration of web site information architecture. *Interacting With Computers*, *20*, 535–544.

Katsanos, C., Tselios, N., & Avouris N. (2010a). Evaluating website navigability: Validation of a tool-based approach through two eye-tracking user studies. *New Review of Hypermedia and Multimedia*, *16*, 195–214.

Katsanos, C., Tselios, N., & Avouris N. (2010b). A survey of tools supporting design and evaluation of websites based on models of human information interaction. *International Journal on Artificial Intelligence Tools*, *19*, 755–781.

Kaviani, N., Finke, M., Fels, S., Lea, R., & Wang, H. (2009). What goes where?: Designing interactive large public display applications for mobile device interaction. In *ICIMCS 2009: Proceedings of the International Conference on Internet Multimedia Computing and Service* (pp. 129–138). New York, NY: ACM Press.

Kim, K., Jacko, J., & Salvendy, G. (2011). Menu design for computers and cell phones: Review and reappraisal. *International Journal of Human–Computer Interaction*, *27*, 383–404.

Kostakos, V., Kukka, H., Goncalves, J., Tselios, N., & Ojala, T. (2013). Multipurpose public displays: How shortcut menus affect usage. *IEEE Computer Graphics and Applications*, *33*, 50–57.

Kostakos, V., & O'Neill, E. (2003). A directional stroke recognition technique for mobile interaction in a pervasive computing world. In *HCI 2003: Proceedings of the 10th International Conference on Human–Computer Interaction* (pp. 197–206). New York, NY: Springer Verlag.

Kurtenbach, G., & Buxton, W. (1993). The limits of expert performance using hierarchical marking menus. In *INTERCHI 1993: Proceedings of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing* (pp. 482–487). New York, NY: ACM Press.

Landauer, T. K., & Dumais S. (2008). Latent semantic analysis. *Scholarpedia*, *3*, 43–56.

Landauer, T. K., McNamara, D. S., Dennis, S., & Kintsch, W. (2007). *Handbook of latent semantic analysis*. New York, NY: Psychology Press.

Larson, K., & Czerwinski, M. (1998). Web page design: implications of memory, structure and scent for information retrieval. In *CHI 1998: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 25–32). New York, NY: ACM Press.

Ling, C., Hwang, W., & Salvendy, G. (2007). A survey of what customers want in a cell phone design. *Behaviour & Information Technology*, *26*, 149–163.

McCarthy, J. F., Costa, T. J., & Liongosari, E. S. (2001). Unicast, outcast & groupcast: Three steps toward ubiquitous, peripheral displays. In *UBICOMP 2001: Ubiquitous Computing, Third International Conference* (pp. 332–345). Berlin, Germany: Springer-Verlag.

Miller, C. S., & Remington, R. W. (2004). Modeling information navigation: Implications for information architecture. *Human–Computer Interaction*, *13*, 225–271.

Miller, D. P. (1981). The depth/breadth tradeoff in hierarchical computer menus. In *HFES 1981: Proceedings of the 25th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 296–300). Thousand Oaks, California: SAGE Publications.

Morville, P. (2005). *Ambient findability: What we find changes who we become*. Sebastopol, CA: O'Reilly Media.

Müller, J., Alt, F., Michelis, D., & Schmidt, A. (2010). Requirements and design space for interactive public displays. In *MULTIMEDIA 2010: Proceedings of the ACM Multimedia Internation Conference* (pp. 1285–1294). New York, NY: ACM Press.

Muller, J., & Kruger, A. (2006). Towards situated public displays as multicast systems. In *UbiqUM 2006: Workshop on Ubiquitous User Modelling, the 17th European Conference on Artificial Intelligence*.

Nacenta, M., Baudisch, P., Benko, H., & Wilson, A. (2009). Separability of spatial manipulations in multi-touch interfaces. In *ACM GI 2009: Proceedings of the Graphics Interface 2009 Conference* (pp. 175–182). Ontario, Canada: Canadian Information Processing Society.

Ojala, T., Kostakos, V., Kukka, H., Heikkinen, T., Linden, T., Jurmu, M., . . . Zanni, D. (2012). Multipurpose interactive public displays in the wild: three years later. *IEEE Computer*, *45*, 42–29.

O'Neill, E., Kaenampornpan, M., Kostakos, V., Warr, A., & Woodgate, D. (2006). Can we do without GUIs? Gesture and speech interaction with a patient information system. *Personal and Ubiquitous Computing*, *10*, 269–283.

Otter, M., & Johnson, H. (2000). Lost in hyperspace: Metrics and mental models. *Interacting with Computers*, *13*, 1–40.

Patsula, P. J., Detenber, B. H., & Theng, Y.-L. (2010). Structure processing of web-based menus. *International Journal of Human–Computer Interaction*, *26*, 675–702.

Pook, S., Lecolinet, E., Vaysseix, G., & Barillot, E. (2000). Control menus: execution and control in a single interactor. In *CHI 2000: Ext. Abstracts of the SIGCHI conference on Human Factors in Computing Systems* (pp. 263–264). New York, NY: ACM Press.

Resnick, M. L., & Sanchez, J. (2004). Effects of organizational scheme and labeling on task performance in product-centered detail web sites. *Human Factors*, *46*, 104–117.

Rooney, C., & Ruddle, R. (2012). Improving window manipulation and content interaction on high-resolution, wall-sized displays. *International Journal of Human–Computer Interaction*, *28*, 423–432.

Sinha, R., & Boutelle, B. (2004). Rapid information architecture prototyping. In *DIS 2004: Proceedings of the Conference on Designing Interactive System* (pp. 349–352). New York, NY: ACM Press.

Smelcer, J. B., & Walker, N. (1993). Transfer of knowledge across computer command menus. *International Journal of Human–Computer Interaction*, *5*, 147–165.

Smith, P. A. (1996). Towards a practical measure of hypertext usability. *Interacting with Computers*, *8*, 365–381.

Toms, E. G. (2002). Information interaction: Providing a framework for information architecture. *Journal of the American Society for Information Science and Technology*, *53*, 855–862.

Tselios, N., Katsanos, C., & Avouris N. (2009). Investigating the effect of hyperlink information scent on users' interaction with a web site. In *INTERACT 2009: Proceedings of the 12th IFIP TC13 Conference on Human–Computer Interaction, LNCS* 5727, *Part II* (pp. 138–142). Berlin, Germany: Springer-Verlag.

Vogel, D., & Balakrishnan, R. (2004). Interactive public ambient displays: transitioning from implicit to explicit, public to personal, interaction with multiple users. In *UIST'04: Proceedings of the 17th annual ACM Symposium on User Interface Software and Technology* (pp. 137–146). New York, NY: ACM Press.

Witten, I. H., & Frank, E. (1999). *Data mining: practical machine learning tools and techniques with java implementations*. Burlington, MA: Morgan Kaufmann.

ABOUT THE AUTHORS

Christos Katsanos is a postdoctoral researcher in the Hellenic Open University's School of Science and Technology. His research interests include human-computer interaction, web accessibility, information architecture and educational technology. Katsanos received a PhD in electrical and computer engineering from the University of Patras. Contact him at E-mail: ckatsanos@eap.gr.

Nikolaos Tselios is an assistant professor of computer science in education in the Department of Educational Sciences and Early Childhood Education, University of Patras. His research interests include educational technology, human–computer interaction, learning analytics and technology acceptance. He's an ACM member. Contact him at E-mail: nitse@ece.upatras.gr.

Jorge Goncalves is a doctoral candidate in the University of Oulu's Department of Computer Science and Engineering. His research interests include social and ubiquitous computing, and the use of technology to motivate participation. Goncalves received an MSc in computer science from the University of Madeira. Contact him at E-mail: jgoncalv@ee.oulu.fi.

Tomi Juntunen is a research assistant in the University of Oulu's Department of Computer Science and Engineering. His research interests include security and human–computer interaction. Juntunen received a BSc in computer science and engineering from the University of Oulu. Contact him at E-mail: tomi.juntunen@ee.oulu.fi.

Vassilis Kostakos is a professor of computer engineering in the University of Oulu's Department of Computer Science and Engineering. His research interests include ubiquitous computing, human–computer interaction, and social systems. Kostakos received a PhD in computer science from the University of Bath. He's an ACM member. Contact him at E-mail: vassilis@ee.oulu.fi.