Map, Diagram, and Web Page Navigation on Mobile Devices: the Effectiveness of Zoomable User Interfaces with Overviews

Stefano Burigat
burigat@dimi.uniud.it

Luca Chittaro
chittaro@dimi.uniud.it

Edoardo Parlato
parlato.edoardo@gmail.com

HCI Lab
Department of Mathematics and Computer Science
University of Udine
Via delle Scienze 206, 33100
Udine, Italy

ABSTRACT

This paper presents a study comparing two Zoomable User Interfaces with Overviews (ZUIOs) against a classic Zoomable User Interface (ZUI) in the context of user navigation of large information spaces on mobile devices. The study aims at exploring (i) if an overview is worth the space it uses as an orientation tool during navigation of an information space and (ii) if part of the lost space can be recovered by switching to a wireframe visualization of the overview and dropping semantic information in it. The study takes into consideration search tasks on three types of information space, namely maps, diagrams, and web pages, that widely differ in structural complexity. Results suggest that overviews bring enough benefit to justify the used space if (i) they highlight relevant semantic information that users can exploit during search and (ii) the structure of the considered information space does not provide appropriate orientation cues.

Categories and Subject Descriptors

H.5.2 [**Information Interfaces and presentation**]: Evaluation, screen design, Graphical user interfaces (GUI); I.3.6 [**Computer Graphics**]: Interaction techniques

General Terms

Experimentation, Human Factors

Keywords

Overview&Detail interfaces, Zoomable User Interfaces, small-screen devices, mobile interaction, experimental evaluation

1. INTRODUCTION

Mobile devices such as PDAs and Smartphones are increasingly being used to navigate large information spaces such as documents, images, web pages, and maps. However, since the common form factors of mobile devices constrain screen space to a small fraction of what is available on a desktop, navigating such information spaces is extremely difficult for mobile users.

In general, navigating a large information space such as a map on a computer screen involves using the screen as a window (hereinafter, *viewport*) onto the larger information space and exploiting appropriate interaction techniques to change the portion of that space that the viewport displays. The most basic approach to do this is to let users move the viewport over the information space. This can be achieved through *scrolling* by means of scrollbars that provide separate horizontal and vertical viewport control or *panning* by directly dragging the information space in any direction. *Zoomable User Interfaces* (ZUIs) [16, 4] combine panning, which has been found to be rather tedious for large information spaces [13], with *zooming*, which can be used to obtain multiple perspectives on the information space by changing its scale [8]. The main drawback of ZUIs is that users tend to lose overview of the information space while navigating [6, 9]. *Overview&Detail* interfaces can reduce the resulting orientation problems. These interfaces are based on displaying an overview of the information space together with a detail view of a portion of that space. The overview is usually a small-scale thumbnail of the whole information space that includes a properly positioned graphical highlight (hereinafter, *viewfinder*) of that portion of space which is currently displayed by the detail view. This is by far the most common configuration, although multiple overviews and multiple detail views can in general be displayed simultaneously [17]. It is also possible to distinguish between overviews that allow users to manipulate the viewfinder to perform panning and possibly zooming of the detail view, and overviews that do not support direct manipulation of the viewfinder and mainly help users orient themselves in the information space. In general, Overview&Detail interfaces offer the following benefits [9]: i) navigation is more efficient if users can navigate by moving the viewfinder within the overview rather than using the detail view [3], ii) the overview aids users in keeping track of their current position in the information space [17], iii) the overview provides users with relevant information for their current task, iv) users feel more in control of the navigation process [21].

Increasingly, ZUIs are combined with Overview&Detail interfaces to provide users with more flexibility in navigating information spaces. Indeed, the zooming capabilities provided by ZUIs are useful in exploring information spaces and overviews

MobileHCI 2008, September 2-5, 2008, Amsterdam, the Netherlands.

Figure 1: Example of Zoomable User Interface with Overviews in a desktop application. An overview is displayed in the bottom right corner and highlights (through a blue viewfinder) which portion of map is displayed by the detail view. Source: Google Maps web site.

mitigate the above mentioned main drawback of ZUIs. *Zoomable User Interfaces with Overviews* (ZUIOs, see Fig. 1) are the result of this combination. ZUIOs are employed and have been shown to be effective for various tasks in the desktop domain [3, 15, 10]. However, they are rarely used in mobile applications. Indeed, while they are an improvement over basic ZUIs and Overview&Detail interfaces, porting a ZUIO to mobile devices exacerbates the typical drawbacks of Overview&Detail interfaces. One major drawback is that the integration of overview and detail views requires mental and motor effort because there is a spatially indirect relation between the two views. Such difficulty in relating the two views might strain memory and increase the time used for visual search of an information space [6], especially when the views are limited by a small screen. An additional drawback, which might be acceptable on desktop systems but becomes very serious on the small screens of mobile devices, is that the overview reduces the space available to the detail view. Moreover, the overview itself must be very small and thus becomes extremely difficult to read, making one doubt of its actual effectiveness [7]. Given these premises, it may seem unsurprising that a recent study [5] reached the conclusion that ZUIOs on mobile devices are ineffective to help users complete search tasks on scatterplots faster than ZUIs. However, the design space for ZUIOs is large and ZUIOs have not been evaluated yet on many information spaces and tasks. Since previous studies on desktop systems showed that different information space and task combinations lead to different results concerning ZUIOs effectiveness, the conclusion reached by Buring et al. [5] should not be overgeneralized.

With this paper, we start an exploration of the interface, information space and task design space for ZUIOs, with the aim of providing useful criteria for practitioners and designers of mobile interfaces. The study we present compares two ZUIO variants with a traditional ZUI in search tasks that involve three types of information spaces, respectively maps, diagrams, and web pages, that differ in structural complexity and information density. We specifically focus on the use of ZUIOs as orientation tools, studying if an overview is worth the space it uses to support user navigation of an information space. Additionally, we aim at determining if the space the overview occupies can be partially recovered by switching to wireframe representation, thus hiding the semantic information the overview provides.

In the following, we first report on prior research on navigation of large information spaces on mobile devices, with a focus on ZUIs and Overview&Detail interfaces. Then, we describe in detail the interfaces we considered in the evaluation and present the experiment we carried out to assess how these interfaces supported users in navigating maps, diagrams, and web pages on a PDA phone. Finally, we discuss the obtained results and present future work.

2. RELATED WORK

In this section, we will discuss relevant literature on ZUIs and Overview&Detail interfaces, describing possible design choices for these interfaces and focusing on studies and applications for mobile devices.

2.1 Zoomable User Interfaces

Since there are multiple ways to actually implement panning and zooming, the design space for ZUIs is large. For example, zoom can be *geometric*, where the scale linearly determines the size of each object in the information space, or *semantic*, where the representation of an object depends on the scale of the information space and objects can change size, shape, details or they can appear/disappear from the visualization when zoomed [16]. The change in scale triggered by a zoom action can be instantaneous or animated (i.e., the transition from the old to the new scale is smooth). Zooming and panning can also be non-linear, such as in goal directed zoom [22], which supports direct zooming to an appropriate scale, Speed-Dependent Automatic Zooming [11], which combines scrolling and zooming into a single operation, and automatic zoom to objects, where selecting an object leads to an automatic zoom on that object. Pad++ [4] is an example of desktop ZUI where panning and zooming are coupled with semantic zooming and animated transitions.

Various ZUIs have been specifically designed for mobile devices. ZoneZoom [18] is a ZUI that lets users easily explore large images on Smartphones: each image is partitioned into nine cells, each one mapped into a number of the phone keypad, and pressing a numeric key produces an automated pan and zoom on the associated cell (which can then be recursively partitioned into nine more cells) while a dedicated key enables users to zoom out. An adaptation of the ZoneZoom technique to PDAs has been proposed by [19] and allows users to pan and zoom on images by interacting with a grid overlaid on the currently displayed image portion. The grid size is proportional to the size of the whole image and each grid cell can be tapped to zoom on the corresponding portion of the image. Cells can also be merged or split to provide users with different zoom levels. Speed-Dependent Automatic Zooming [11] has been adapted to mobile devices by Jones et al. [12]. In the mobile version, two concentric circles are drawn when users tap on the information space with the pointer. If the pointer remains within the inner circle, the user is free to pan in any direction and the panning rate increases as the pointer moves away from the starting position. When the pointer moves beyond the inner circle, both zooming and panning operations take place. The information space progressively zooms out as the user moves closer to the outer circle and the panning speed changes to maintain a consistent visual flow. When the pointer reaches the outer circle, no further zooming occurs, while panning remains active. The experimental evaluation [12] showed that the proposed technique reduces the physical navigational workload of users with respect to a standard interface based on the use of scrollbars, panning and zoom buttons.

2.2 Overview&Detail interfaces

A number of studies have investigated the usability of Overview&Detail interfaces on desktop computers and have found that they improve user satisfaction and efficiency over detail-only interfaces. Beard and Walker [3] investigated the effect of providing an overview compared to having only scrollbars. Subjects used an overview that allowed dragging the viewfinder and one that allowed both drag and resize of the viewfinder. For tasks where subjects tried to locate a word in a tree and tasks where they repeatedly went from one side of the tree to the other, the Overview&Detail interface lead to significantly faster task completion. North and Shneiderman [15] compared a detail-only, a classic Overview&Detail interface, and an Overview&Detail interface with no viewfinder within the overview. Compared to the detail-only interface, the classic Overview&Detail interface was faster and scored significantly higher on a satisfaction questionnaire. Hornbaek and Frokjaer [10] compared the usability of a detail-only, a fisheye, and an Overview&Detail interface for electronic documents. They found that essays produced with the aid of the Overview&Detail interface were scored significantly higher than essays produced with the aid of the detail-only interface. However, for tasks that required subjects to read and answer a specific question, the Overview&Detail interface was slower compared to the detail-only interface. All but one of the subjects preferred the Overview&Detail interface. Baudisch et al. [2] analyzed the same interfaces in web browsing tasks and found that user performance with the fisheye interface was superior than performance with the other two interfaces, but users preferred the Overview&Detail interface nonetheless. Hornbaek et al. [9] compared the usability of ZUIs with and without overviews. Subjects solved browsing and navigation tasks on two maps, where one map was organized in multiple levels and allowed for semantic zooming. Subjects were faster without the overview when using the map with multiple levels. The authors assumed that this type of information space provides richer navigation cues and thus makes an overview unnecessary. In an evaluation carried out by Nekrasovski et al. [14], overviews were combined with a ZUI and a Focus&Context interface for navigating large hierarchical trees but were not found to improve performance, albeit they were perceived to be beneficial by users.

While research has extensively investigated Overview&Detail interfaces on desktop systems, to the best of our knowledge only two recent studies investigated these interfaces on mobile devices. In the study by Buring et al. [5], participants performed search tasks on scatterplots by using two interfaces on a PDA, one displaying a detail view and an overview and the other displaying only the detail view. In the first interface, the detail view was much smaller than in the second interface because the overview did not overlap it. While there was no significant difference in user preference between the interfaces, participants solved search tasks faster without the overview. This may indicate that, on small screens, a larger detail view can outweigh the benefits gained from an overview window. Results also revealed that individual differences in spatial ability did not have a significant effect on task completion times. Roto et al. [20] compared instead two different methods to visualize web pages. The first method is based on reformatting web pages into a single column that fits the screen width. The second method, called Minimap, combines a restructuring approach that changes the size of text and limits the maximum width of text paragraphs to the width of the browser viewport, with an overview of the web page that is overlaid on top of the browser viewport and helps users to locate information inside the page. Their study shows that Minimap scores better in usability ratings and is preferred by users. Unfortunately, since Minimap is a combination of two approaches, it is impossible to determine the contribution of the overview to the results. While the authors claim that the pros of the overview overshadowed the cons, they also provide preliminary results of an additional study that found the Minimap to work better than the single column approach even without using the page overview.

3. THE CONSIDERED INTERFACES

Figure 2 shows the three interfaces we considered in our study. The first interface (hereinafter, *Classic ZUI*) is a ZUI implementation that is widely employed in desktop as well as mobile applications (Fig. 2a). It allows users to perform panning by dragging the portion of information space displayed in the viewport. On PDAs, dragging is carried out by moving the stylus in any direction while keeping it in contact with the screen. While the stylus moves, the viewport moves the same distance in the opposite direction so that the information space moves in the direction indicated by the user. Users can zoom by choosing a specific zoom level among a predefined set. The current level can be changed by tapping on two icons, respectively depicting a plus and a minus sign, that are located at the right and left corner in the upper area of the screen. Both icons are overlaid on the information space but are transparent to minimize occlusion. The plus (minus) icon highlights when tapped to indicate that it is in its pressed state and grays out when the user reaches the maximum (minimum) zoom level.

The second interface (hereinafter, *Classic ZUIO*) combines an Overview&Detail implementation (without direct manipulation of the viewfinder) with the panning and zooming features of the Classic ZUI interface. The overview of the considered information space is displayed as a small thumbnail at the bottom right corner of the detail view and contains a viewfinder that highlights which is the portion of space displayed in the detail view (Fig. 2b).

The third interface (hereinafter, *Wireframe ZUIO*) differs from the Classic ZUIO in the way it displays the overview and the viewfinder. Indeed, both the overview and the inner viewfinder are simple rectangular outlines overlaid on the information space at the bottom right corner of the detail view and do not display any semantic content concerning the information space (Fig. 2c).

The main difference among the three interfaces is that they provide users with a different amount of orientation cues. As previously discussed, overviews can in general act as a panning and zooming tool by allowing direct manipulation of the viewfinder or as an orientation tool that aids users in keeping track of their current position in the information space. In our study, we focused on this latter role, hence both ZUIOs display a non-manipulable overview. Classic ZUI does not provide additional orientation cues other than those contained in the information space (e.g., landmarks on maps), while Classic ZUIO provides additional orientation cues through the viewfinder position and the semantic content in the overview. Wireframe ZUIO is an intermediate solution that keeps the information provided by the viewfinder but drops the semantic content in the overview to reduce the amount of space it occludes on the detail view. We have chosen these interface designs to investigate what is the best tradeoff between space used by an overview and orientation cues the overview provides to users during navigation. The three interfaces employ the same panning and zooming technique, based on dragging the portion of information space displayed in the detail view and

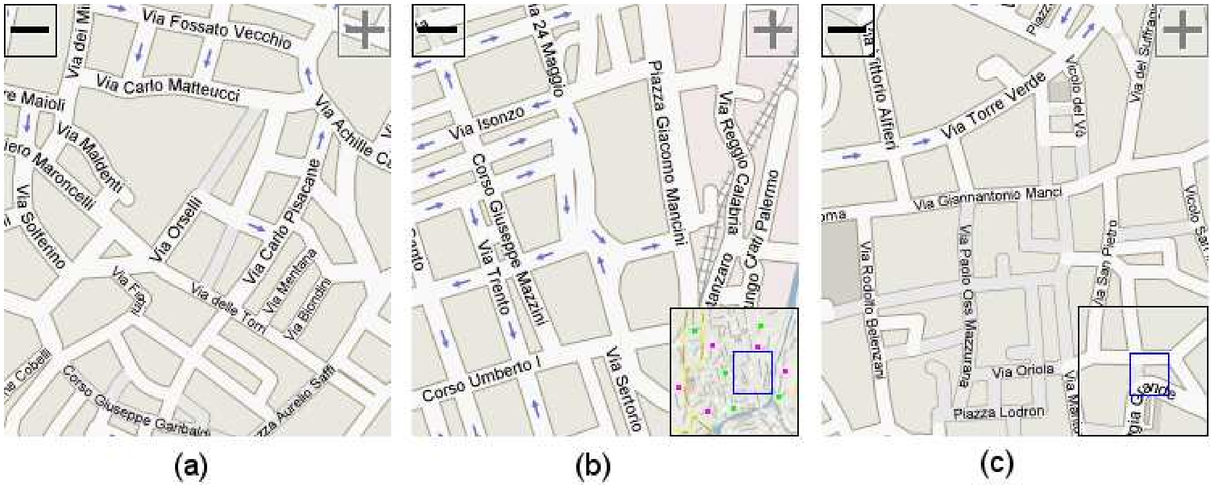

Figure 2: The three interfaces we considered in the evaluation: Classic ZUI (a), Classic ZUIO (b), Wireframe ZUIO (c). In all interfaces, panning is carried out by dragging the information space; zooming is carried out by tapping on the two icons depicting a plus and a minus sign. In both ZUIOs, the overview is displayed in the lower right corner, is tightly coupled with the detail view, and the viewfinder within the overview cannot be manipulated. In Wireframe ZUIO, the overview is displayed as a wireframe and does not contain semantic information about the information space.

using zoom icons. We chose this specific technique because of its widespread use in the mobile domain, which should make our results particularly interesting for practitioners who need to evaluate the benefits of overviews for their applications.

The two ZUIOs have three other common features that have an impact on their effectiveness. The first feature is that overview and detail views are *tightly coupled* [1], i.e., manipulation of the detail view (panning, zooming) is immediately reflected in the overview as variations in the position or size of the viewfinder during navigation. As the results of several studies show (e.g. [15]), this is indeed an essential feature for navigation support. The second common feature is the *zoom factor*, i.e., the level of magnification between the two views. Studies in the desktop domain have shown that the zoom factor should have a value lower than 25-30 [17, 21], and Plaisant et al. [17] recommend to add intermediate overview levels, as needed, if the zoom factor is above these thresholds. In our case, the maximum zoom factor users could reach during the evaluation was 10, i.e., the detail view was at most ten times larger than the overview, while the minimum zoom factor was slightly larger than 3. The third feature is the *size of the overview*, which influences the zoom factor and determines the amount of information that users can see through the overview. A large overview should thus be preferable, but the larger the overview the smaller the screen real estate available to the detail window. In general, Plaisant et al. [17] argue that the size of the overview and the detail view are task dependent. Moreover, the proportions of the overview typically depend on the proportions of the considered information space, to avoid introducing distortions. In our case, the overview fills an 80x80 pixels area on a 240x320 screen (i.e., one third of the screen horizontally and one fourth vertically), since all our information spaces were square. We informally determined the overview size by presenting different configurations to a small sample of users (distinct from the ones participating in the main evaluation) and asking them to select their preferred configuration. Users found smaller overviews to be too difficult to read while they found that larger overviews occluded the detail view too much.

4. EXPERIMENTAL EVALUATION

The experimental evaluation we carried out to compare the interfaces described in the previous section required users to navigate maps, diagrams, and web pages on a mobile device, searching for specific targets. We chose these three information spaces for various reasons. First, no previous investigation on navigating them with mobile ZUIOs appears in the literature (as previously discussed, there is no way to determine what was the effect of the overview in the Minimap study [20]). Second, mobile applications and services based on these information spaces are becoming increasingly common, so that investigating how to best support users in navigating them can be immediately useful to designers and practitioners. Third, the differences among the three information spaces make it possible for us to study the effects of structure and density (i.e., the number of objects) of an information space on the effectiveness of ZUIOs.

4.1 Participants

The evaluation involved a sample of 24 users (13 male, 11 female). They were undergraduate students from diverse backgrounds (9 Computer Science, 3 Humanities, 2 Economics, 2 Engineering) or people from other occupations (3 secondary school students, 2 clerks, 2 consultants, 1 teacher). Their age ranged from 16 to 37, averaging at 26. Almost half of them (10 out of 24) had occasionally used a PDA before. Most of them had used maps, diagrams and web pages on desktop PCs (respectively, 18, 18 and 23 users out of 24). Four users had occasionally used maps on PDAs and 7 users had occasionally used web pages on PDAs. Users were volunteers recruited by direct or email contact.

4.2 Materials

During the experiment, participants were provided with a 200Mhz Windows Mobile 5.0 PDA phone with a 2.8" display and QVGA (240x320) resolution. Figure 3 shows the Classic ZUIO on the PDA phone. With all interfaces and information spaces we considered, the detail view covers a 240x268 area in the middle, while two menu bars display at the top and at the bottom of the screen.

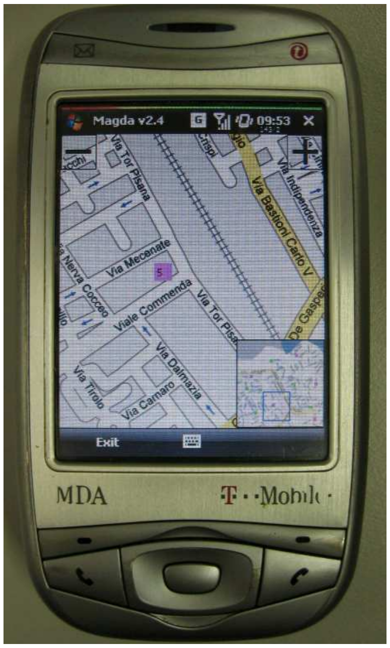

Figure 3: The PDA phone we used for the evaluation.

Figure 4 shows examples of the information spaces used in the experiment: city maps, diagrams and web pages. Maps depicted a typical city area and included 10 numbered color icons that were used as possible targets for the evaluation. Diagrams were made of connected color rectangles with a descriptive text. Web pages were taken from news sites and had a central area with short titled sections and two side areas with links. Maps are an example of unstructured information space with high information density, where a high number of objects are mixed together. Web pages have high information density but a well-defined structure that simplifies navigation when users are searching for specific content. Diagrams have low information density and a well-defined structure. The structure of each information space was visible at all zoom levels although finer details such as street names, diagram texts or web page words were not readable at the coarsest levels. The size of all information spaces at the maximum zoom level was 800x800 pixels. All information spaces were initially presented at the lowest zoom level so that they were almost entirely visible at the beginning of tasks. Four zoom levels were available to users, thus requiring three taps on the zoom-in icon to move from the lowest to the highest zoom level.

4.3 Tasks

Each participant was assigned three *navigation* tasks and three *spatial memory acquisition* tasks for each interface (18 tasks in total). Each one of the three navigation tasks was carried out on a different information space, one task on a map (MapTask), one on a diagram (DiagramTask), and one on a web page (WebTask). A spatial memory acquisition task (SpatialMemoryTask) followed each navigation task. The aim of the SpatialMemoryTask was to assess which interface enabled users to develop a better mental map of the information space. Tasks are described in detail in the following:

- The MapTask required users to navigate a city map to find the location of two specific targets. In particular, users were informed that they had to search for two specific hotels and that all hotels were depicted as numbered color icons (as is commonly done in tourist applications such as search engines for points of interest). An example of this task is: “Find out where hotels 1 and 3 are located on the map and tap on their icons as soon as you locate them”. The two hotels to search for were located in opposite areas of the map so that users had to perform a thorough exploration to find them. In the Classic ZUIO, hotels were clearly visible in the overview as small color dots.

- The DiagramTask required users to navigate a diagram to find two specific nodes. An example of this task is: “Find out the location of the ‘Graduate students’ node and the ‘Support staff’ node, and tap on them as soon as you locate them”. As with maps, the two nodes to search for were located in opposite areas of the diagram.

- The WebTask required users to navigate a web page to identify the occurrence of two specific words in the text. An example of this task is: “Find one occurrence of the word ‘radio’ and one occurence of the word ‘school’ in the text sections of the web page and tap on each of them”. As with maps and diagrams, targets were located in opposite areas of the page.

- The SpatialMemoryTask required users to mark the exact location of the two targets previously searched for on a paper reproduction of the considered information space at the lowest zoom level (i.e., at the zoom level where less detail was available in the information space). An example of this task is: “Point out where the hotels you have just searched for are located in this paper map”. To perform the SpatialMemoryTask, participants could only rely on spatial knowledge previously acquired during navigation, since no zooming or panning was allowed during this task to read details on the considered information space.

The choice of the tasks for the study was mainly driven by the realization that, in many mobile scenarios, users need to quickly search an information space or need to retrieve information they have previously searched for to support their current activities. Many of these activities are geographic in nature, such as locating certain points of interest, and are thus based on maps. A common feature of map-based applications (e.g., navigation systems, mobile guides, GIS) is that they allow users to highlight the objects of interest that can be relevant to their activities. Hence, we designed the map search task to reflect this common scenario. Since highlighting is such an essential feature in map-based applications, we applied it to the overview in the Classic ZUIO as well. Without highlighting, the overview provides only basic orientation information, which is exactly what the Wireframe ZUIO does. To obtain results that could be compared among different information spaces we designed the tasks for web pages and diagrams to be similar to the ones for maps. However, with those information spaces there was no need of highlighting targets since they were actually visible in the overview and corresponded to well defined objects (text sections in the first case, text within blocks in the second).

4.4 Procedure

The experimental design was within-subjects. Participants were initially briefed by the experimenter about the nature of the experiment. Then, they were provided with an introduction and demonstration of the interfaces. This was followed by the

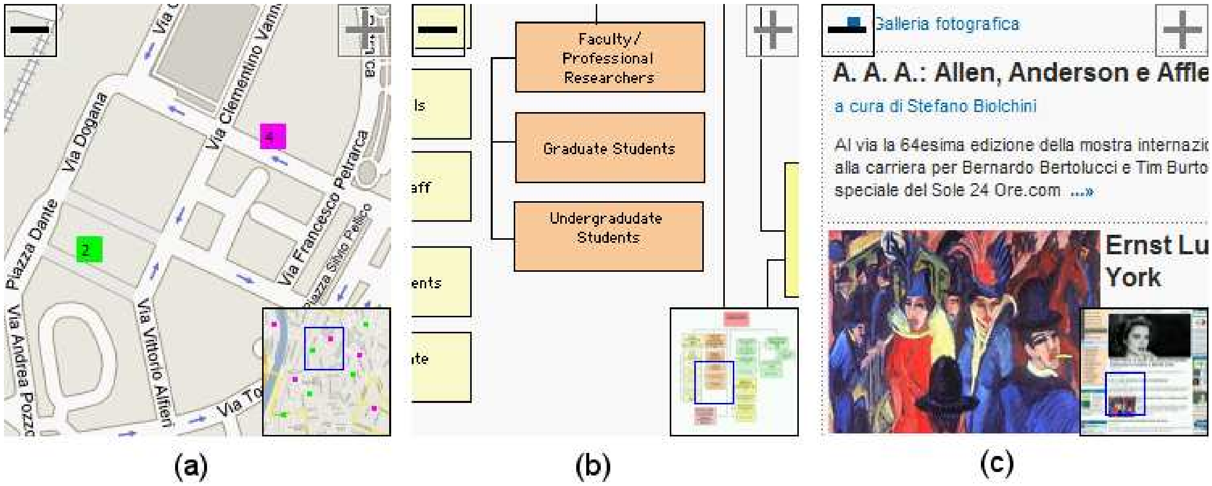

Figure 4: Examples of the three information spaces we considered in the evaluation: maps (a) diagrams (b), web pages (c).

presentation of training tasks for one of the information spaces to let users familiarize with the interfaces. Training tasks were similar to the experimental tasks but involved different targets. While carrying out the training tasks, participants were allowed to ask any question to the experimenter to clarify possible doubts concerning the interfaces or the tasks to be carried out. After training, users carried out the experimental tasks for the considered information space. Subsequently, users were asked to express their preferences by ordering the three interfaces from the best one to the worst one according to their ease of use and the usefulness of the provided information. The same procedure was then followed for the other information spaces. At the end of the experiment, participants were briefly interviewed to collect their comments.

Printed task sheets (containing the description of each task) were presented by the experimenter to participants one at a time, to provide clear instructions for each task and easy reference to target names during task execution. For each task, participants were asked first to read the task description on the printed sheet. Except for the SpatialMemoryTask, users (when ready) were then required to tap on a "Start Task" button that was initially displayed by the PDA and perform the task. Each task ended when users tapped on the last target. The SpatialMemoryTask did not require users to interact with the PDA phone and ended when users marked the last target on the paper reproduction of the considered information space. All possible care was taken to counterbalance learning effects due to repetitive testing:

- Every user was presented with a different order of the information spaces.
- For each information space, every user was presented with a different order of the experimental conditions.
- During testing, target positions were different for each test condition. Four configurations of targets were produced for every information space, one for the training phase and three for the testing phase. Configurations were kept as similar as possible in terms of complexity, i.e. relative distance of targets.
- There was no fixed association between test condition and target configuration. This way, a condition could not benefit by possibly unaccounted factors that might make a target configuration easier to complete than others.

4.5 Logged data

Logging code automatically recorded the following data for each task carried out by each participant:

- Task completion time: the time taken to complete the task, defined as the time elapsed between the tap on the "Start Task" button and the tap on the last target.
- User interface actions: the number of distinct user interface actions to carry out the task. A pan action was recorded when users dragged the stylus on the information space. A zoom action was recorded when users tapped on zoom buttons. A target selection action was recorded when users tapped on targets on the screen.
- Pan timings: the start time, end time and duration of each pan action. A pan action starts when users start dragging the stylus on the information space and ends when users lift the stylus from the information space.
- Accuracy: the distance of the actual target location from the location indicated by the user in the SpatialMemoryTask.

4.6 Hypotheses

Since all the considered tasks ultimately involve the exploration of information spaces at a fine level of detail (i.e., a high zoom factor), the advantages of overviews in terms of orientation cues should overcome the disadvantage of having part of the detail view hidden (because high zoom factors imply that the overview covers only a small fraction of the whole information space). Thus, our hypotheses in the present study are the following:

- Since the overview displays possible targets, users should be faster with Classic ZUIO than they are with Wireframe ZUIO and Classic ZUI in all search tasks.
- Users should be faster in carrying out search tasks with Wireframe ZUIO than they are with Classic ZUI, because (although limited) more orientation cues are provided.
- Users should be more accurate in identifying the location of targets in the SpatialMemoryTask with both ZUIOs, because of the orientation cues provided by these interfaces.

5. RESULTS

In this section, we describe our findings in detail.

5.1 Task completion times

Figure 5 shows the mean completion times for map, diagram and web page tasks. Task completion times were subjected to the Kolmogorov-Smirnov test of normality prior to further analysis. The test did not reveal deviations from the normal distribution. A one-way analysis of variance (ANOVA) was thus employed on the times. The within-subjects factor was the type of interface with three levels: Classic ZUI (ZUI in tables), Wireframe ZUIO (WZUIO in tables), Classic ZUIO (CZUIO in tables). The ANOVA revealed a significant effect only for maps ($F(2, 69) = 3.325, p < 0.05$). Therefore, we employed the Tukey post-hoc test for comparison among pairs of means, which showed that users spent significantly less time to search for targets with Classic ZUIO than they did with Classic ZUI ($q = 3.581, p < 0.05$).

5.2 User interface actions

We employed Friedman's test to analyze the number of zoom and pan actions performed by users to complete each task. Means are shown in Fig. 6 and 7. For zoom actions, the analysis revealed a significant effect in MapTasks ($T = 13.14, p < 0.01$) as well as DiagramTasks ($T = 10.86, p < 0.005$). Dunn's Multiple Comparison post-hoc test was then used to compare pairs of means. A statistically significant difference was found in the number of zoom actions between Classic ZUIO and Classic ZUI ($p < 0.01$), with Classic ZUIO requiring fewer actions in both MapTasks and DiagramTasks. For pan actions, the analysis revealed a significant effect in MapTasks ($T = 8.579, p < 0.05$) and the post-hoc test pointed out a statistically significant difference in the number of pan actions between Classic ZUIO and Classic ZUI ($p < 0.05$), with Classic ZUIO requiring less actions.

5.3 Pan time

As with task completion times, we used the Kolmogorov-Smirnov test of normality prior to further analysis and the test did not reveal deviations from the normal distribution. ANOVA was then used to analyze pan times, whose means are shown in Fig. 8. The ANOVA did not reveal any statistical significant difference among mean times for any information space.

5.4 Errors in the SpatialMemoryTask

Friedman's test was used to analyze the number of errors made by users in the SpatialMemoryTask (Fig. 9), where an error was counted when the Euclidean distance between the target location estimated by the user, and the actual location of the target on the information space was higher than a predefined threshold (corresponding to 70 pixels at the maximum 800x800 resolution). This threshold marked a likely area of approximation for an accurate answer, since beyond this area it was possible to find other (wrong) targets. The analysis pointed out a significant effect in the SpatialMemoryTask for maps ($T = 9.781, p <$

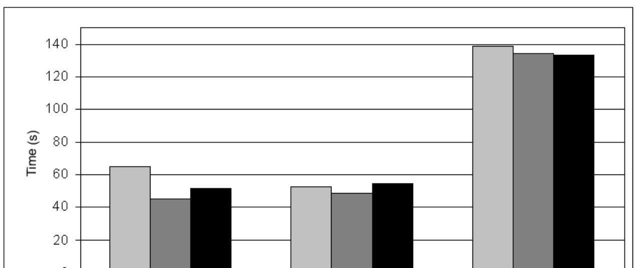

	Maps	Diagrams	Web pages
ZUI	64,57	52,54	138,9
CZUIO	44,99	48,71	134,2
WZUIO	51,52	54,59	133,4

Figure 5: Mean completion times for map, diagram and web page tasks.

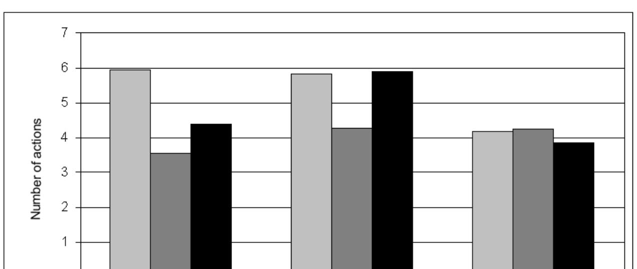

	Maps	Diagrams	Web pages
ZUI	5,952	5,815	4,19
CZUIO	3,571	4,259	4,238
WZUIO	4,381	5,905	3,857

Figure 6: Mean number of zoom actions for map, diagram and web page tasks.

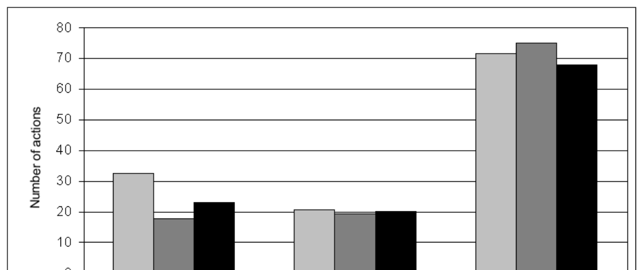

	Maps	Diagrams	Web pages
ZUI	32,7	20,67	71,76
CZUIO	17,73	19,38	75,1
WZUIO	23,13	20,33	68,1

Figure 7: Mean number of pan actions for map, diagram and web page tasks.

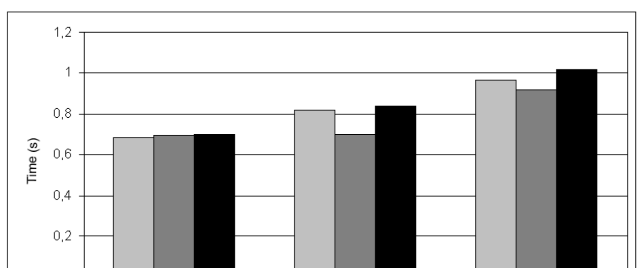

	Maps	Diagrams	Web pages
ZUI	0,686	0,823	0,967
CZUIO	0,696	0,701	0,919
WZUIO	0,703	0,838	1,019

Figure 8: Mean pan times for map, diagram and web page tasks.

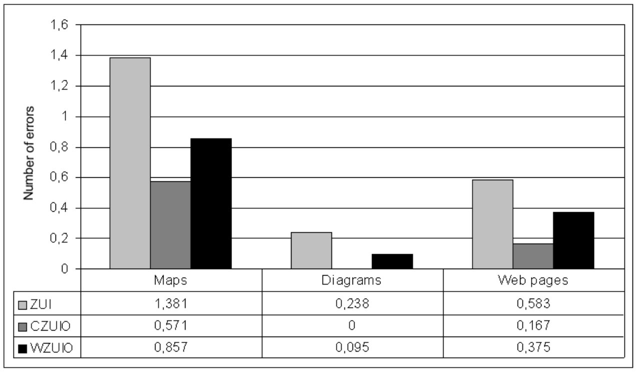

Figure 9: Mean number of errors in the SpatialMemoryTask for maps, diagrams and web pages.

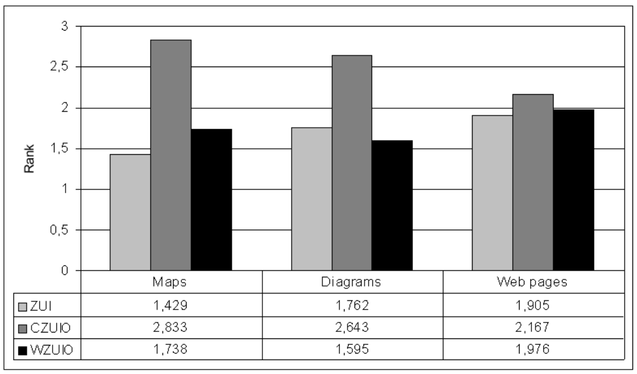

Figure 10: Mean preference for each interface (higher numbers correspond to better scores).

0.01), diagrams ($T = 6.333, p < 0.05$), and web pages ($T = 7.682, p < 0.05$). The subsequent post-hoc test (Dunn's Multiple Comparisons) revealed a statistically significant difference between Classic ZUIO and Classic ZUI ($p < 0.05$) for maps, but was not able to identify significant differences between pair of conditions for diagrams and web pages.

5.5 Subjective preference

To analyze the data on subjective preference, we employed Friedman's test. Since users were asked to rate the three interfaces from the best to the worst, we assigned a score of 3, 2, 1 respectively to the first, second, and third interface. An appropriate fractionary score was assigned to draws, which were allowed. The analysis pointed out a significant effect for maps ($T = 24.96, p < 0.0001$) as well as diagrams ($T = 15.75, p < 0.001$). For maps, Dunn's test for post-hoc analysis among total ranks revealed a statistically significant difference in preference between Classic ZUIO and Classic ZUI ($p < 0.001$) as well as between Classic ZUIO and Wireframe ZUIO ($p < 0.01$), with users preferring the first interface in both cases. The same result was obtained for diagrams, with user preference for Classic ZUIO being significantly higher than preference for Classic ZUI ($p < 0.01$) and Wireframe ZUIO ($p < 0.05$).

6. DISCUSSION

Overall, the experiment partially confirmed two out of the three hypotheses we made. Classic ZUIO outperformed Classic ZUI in terms of time needed by users to complete MapTasks, but no statistically significant difference was found in DiagramTasks and WebTasks. This result suggests that the effectiveness of ZUIOs for map search tasks is higher when the overview displays semantic content and highlights the objects of interest for users, as was the case of hotels in MapTasks. Without overview, users had to fully explore the considered map to search for targets. With the overview, users could instead immediately identify those map areas that contained targets and directly navigate to them.

The effectiveness of Classic ZUIO with maps was also confirmed by the analysis of user interface actions. Indeed, users performed significantly less zoom and pan actions with this interface than they did with Classic ZUI. While the result for pan actions is strictly related to the result for task completion times, the result for zoom actions is indicative of a different navigation strategy adopted by users. When the overview was available, users were more prone to zoom in up to the maximum level of detail and then navigate at that level using the overview as guidance, while they typically went through a higher number of zoom actions with the other two interfaces, zooming out as they felt the need to get an overview of the explored information space. This behavior extended also to search tasks on Diagrams, where users performed significantly less zoom actions with Classic ZUIO than they did with Classic ZUI. However, this did not have an effect on the time required to complete DiagramTasks.

The fact that Classic ZUIO was as effective as the other two interfaces with diagrams and web pages was probably due to the structure of these information spaces, which provides intrinsic orientation cues to support user navigation. Indeed, users could easily follow the arcs connecting nodes to thoroughly explore diagrams avoiding empty areas of the information space, while the well defined layout of web pages helped users in easily identifying the areas where text sections were located. Thus, in these information spaces, users did not need to use the overview as much as with maps to complete their assigned tasks.

We did not find Wireframe ZUIO to provide any clear benefit over Classic ZUI in any navigation task. Indeed, it could not ouperform Classic ZUI in MapTasks (likely because it lacks the visualization of semantic content in the overview) while it did not make any difference in diagrams and web page navigation because of the structure of these information spaces, as mentioned above for Classic ZUIO.

Both ZUIOs allowed users to be more accurate in the SpatialMemoryTask with respect to Classic ZUI. However, while there was globally a significant effect in all three information spaces, only the difference betweeen Classic ZUIO and Classic ZUI in maps reached significance in the post-hoc analysis. Again, the structure of diagrams and web pages could have helped users in being more accurate with these two information spaces with all three interfaces, thus reducing the advantage provided by the overview.

Finally, preference analysis highlighted a clear user preference for Classic ZUIO over the other two interfaces for both map and diagram navigation. In the first case, this is consistent with performance results, while for diagram navigation this result highlights that users preferred having an overview even if it did not provide significant benefits in terms of navigation effectiveness. No interface was instead clearly preferred for web page navigation but user comments highlighted that the overview was considered detrimental to web navigation because it overlapped the detail view. This drawback was not pointed out by users for the other two information spaces.

7. CONCLUSIONS AND FUTURE WORK

Our study highlights that ZUIOs effectiveness in search tasks is highly dependent on the type of information overviews can provide and on the structure of the considered information space. Overviews that do not provide semantic information, such as in the case of the Wireframe ZUIO, have only a limited effect on user navigation performance and are negatively judged by users, despite their reduced screen occupation. Overviews that highlight semantic information can instead improve user navigation performance and are quite effective as orientation tools, especially when the information space does not have a well defined structure that provides orientation cues during navigation. Moreover, ZUIOs seem to be useful when it is important to quickly determine where the user is located with respect to the whole information space or where different objects are located with respect to each other. Thus, examples of applications that would benefit from the use of ZUIOs are mobile tourist guides where maps are used to display the location of points of interest (e.g., monuments, metro stations, hotels), mobile games involving maps with dynamic content (e.g., real-time strategy games), emergency management applications where maps display the location of the members of field teams, equipment and critical events.

However, the results of our study should not be overgeneralized since they concern specific combinations of interface, information space and task. Further evaluations are needed to assess ZUIOs effectiveness under other conditions so that a more complete framework for the use of overviews on mobile devices could be developed. More specifically, an important next step is to explore the use of overviews as panning and zooming tools that allow users to directly navigate an information space by manipulating the viewfinder. In particular, the time needed by users to pan an information space could be significantly reduced by moving the viewfinder within the overview, compared to dragging the information space itself in the detail view. However, interactive overviews may suffer from scalability issues when the zoom factor is high since, in this case, the viewfinder may shrink too much in size and make its manipulation more difficult for users. Scalability issues are directly related to the size of overviews which is a fundamental parameter of ZUIOs, as we discussed in previous sections. In this paper, we set the size of the overview to a fixed value and used only information spaces that could fit exactly that size. Future studies will analyze the effect of size and proportions of the overview, considering that, in general, they should depend on the size and proportions of the considered information space.

8. ACKNOWLEDGMENTS

The authors acknowledge the financial support of the Italian Ministry of Education, University and Research (MIUR) within the FIRB project number RBIN04M8S8.

9. REFERENCES

[1] C. Ahlberg and B. Shneiderman. Visual information seeking: Tight coupling of dynamic query filters with starfield displays. In *Proc. Conference on Human Factors in Computing Systems (CHI 94)*, pages 313–317. ACM Press, 1994.

[2] P. Baudisch, B. Lee, and L. Hanna. Fishnet, a fisheye web browser with search term popouts: a comparative evaluation with overview and linear view. In *Proc. Conference on Advanced visual interfaces (AVI 2004)*, pages 133–140. ACM Press, 2004.

[3] D.B. Beard and J.Q. Walker. Navigational techniques to improve the display of large two-dimensional spaces. *Behaviour and Information Technology*, 9(6):451–466, 1990.

[4] B. Bederson and J. Hollan. Pad++: a Zooming Graphical Interface for Exploring Alternate Interface Physics. In *Proc. Conference on User Interface Software and Technology (UIST 1994)*, pages 17–26. ACM Press, 1994.

[5] T. Buring, J. Gerken, and H. Reiterer. Usability of Overview-supported Zooming on Small Screens with Regard to Individual Differences in Spatial Ability. In *Proc. Conference on Advanced Visual Interfaces (AVI 2006)*, pages 233–240. ACM Press, 2006.

[6] S.K. Card, J.D. Mackinlay, and B. Shneiderman. *Readings in Information Visualization*. Morgan Kaufmann, 1999.

[7] L. Chittaro. Visualizing Information on Mobile Devices. *IEEE Computer*, 39(3):40–45, 2006.

[8] C. Gutwin and C. Fedak. Interacting with Big Interfaces on Small Screens: a Comparison of Fisheye, Zoom, and Panning Techniques. In *Proc. Conference on Graphics Interface (GI 2004)*, pages 145–152. Canadian Human-Computer Communications Society, 2004.

[9] K. Hornbaek, B.B. Bederson, and C. Plaisant. Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM Transactions on Computer-Human Interaction*, 9(4):362–389, 2002.

[10] K. Hornbaek and E. Frokjaer. Reading of Electronic Documents: The Usability of Linear, Fisheye, and Overview+Detail Interfaces. In *Proc. Conference on Human Factors in Computing Systems (CHI 2001)*, pages 293–300. ACM Press, 2001.

[11] T. Igarashi and K. Hinckley. Speed-dependent Automatic Zooming for Browsing Large Documents. In *Proc. Conference on User Interface Software and Technology (UIST 2000)*, pages 139–148. ACM Press, 2000.

[12] S. Jones, M. Jones, G. Marsden, D. Patel, and A. Cockburn. An Evaluation of Integrated Zooming and Scrolling on Small Screens. *International Journal of Human-Computer Studies*, 63(3):271–303, 2005.

[13] V. Kaptelinin. A Comparison of Four Navigation Techniques in a 2D Browsing Task. In *Proc. Conference on Human Factors in Computing Systems (CHI 1995)*, pages 282–283. ACM Press, 1995.

[14] D. Nekrasovski, A. Bodnar, J. McGrenere, F. Guimbretiere, and T. Munzner. An Evaluation of Pan&zoom and Rubber Sheet Navigation with and without an Overview. In *Proc. Conference on Human Factors in Computing Systems (CHI 2006)*, pages 11–20. ACM Press, 2006.

[15] C. North and B. Shneiderman. Snap-together visualization: can users construct and operate coordinated visualizations? *International Journal of Human-Computer Studies*, 53(5):715–739, 2000.

[16] K. Perlin and D. Fox. Pad: an alternative approach to the computer interface. In *Proc. Conference on Computer Graphics and Interactive Techniques (SIGGRAPH 1993)*, pages 57–64. ACM Press, 1993.

[17] C. Plaisant, D. Carr, and B. Shneiderman. Image-Browser Taxonomy and Guidelines for Designers. *IEEE Software*, 12(2):21–32, 1995.

[18] D.C. Robbins, E. Cutrell, R. Sarin, and E. Horvitz. ZoneZoom: Map Navigation for Smartphones with Recursive View Segmentation. In *Proc. International Conference on Advanced visual interfaces (AVI 2004)*, pages

231–234. ACM Press, 2004.

[19] R. Rosenbaum and H. Schumann. Grid-based Interaction for Effective Image Browsing on Mobile Devices. In *Proc. of the SPIE, Volume 5684*, pages 170–180. SPIE - The International Society for Optical Engineering, 2005.

[20] V. Roto, A. Popescu, A. Koivisto, and E. Vartiainen. Minimap: a Web Page Visualization Method for Mobile Phones. In *Proc. Conference on Human Factors in Computing Systems (CHI 2006)*, pages 35–44. ACM Press, 2006.

[21] B. Shneiderman. *Designing the User Interface*. Addison-Wesley, 1998.

[22] A. Woodruff, J. Landay, and M. Stonebreaker. Goal-directed zoom. In *Proc. Conference on Human Factors in Computing Systems (CHI 98)*, pages 305–306. ACM Press, 1998.