

Visualizing references to off-screen content on mobile devices: A comparison of Arrows, Wedge, and Overview + Detail

Stefano Burigat*, Luca Chittaro

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

ARTICLE INFO

Article history:

Received 21 July 2010

Received in revised form 22 December 2010

Accepted 11 February 2011

Available online 26 February 2011

Keywords:

Visualization

Mobile devices

Contextual Cues

Peripheral awareness

Off-screen objects

Overview + Detail

ABSTRACT

When navigating large information spaces on mobile devices, the small size of the display often causes relevant content to shift off-screen, greatly increasing the difficulty of spatial tasks such as planning routes or finding points of interest on a map. Two possible approaches to mitigate the problem are Contextual Cues, i.e. visualizing abstract shapes in the border region of the view area to function as visual references to off-screen objects of interest, and Overview + Detail, i.e., simultaneously displaying a detail view and a small-scale overview of the information space. In this paper, we compare the effectiveness of two different Contextual Cues techniques, Wedge (Gustafson et al., 2008) and Scaled Arrows (Burigat et al., 2006), and a classical Overview + Detail visualization that highlights the location of objects of interest in the overview. The study involved different spatial tasks and investigated the scalability of the considered visualizations, testing them with two different numbers of off-screen objects. Results were multifaceted. With simple spatial tasks, no differences emerged among the visualizations. With more complex spatial tasks, Wedge had advantages when the task required to order off-screen objects with respect to their distance from the display window, while Overview + Detail was the best solution when users needed to find those off-screen objects that were closest to each other. Finally, we found that even a small increase in the number of off-screen objects negatively affected user performance in terms of accuracy, especially in the case of Scaled Arrows, while it had a negligible effect in terms of task completion times.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Today, the capabilities of mobile devices make it possible to navigate large information spaces such as maps to carry out spatial tasks like planning routes, looking for suitable points of interest in a specific area, or viewing the real-time location of individual first responders during emergencies. Unfortunately, the small screen of mobile devices greatly increases the complexity of these activities compared to desktop systems (Chittaro, 2006). Typically, when the information space is displayed in its entirety, users obtain an overview without sufficient detail (e.g., they are unable to read text). By zooming-in, users may obtain needed details but have no direct visual access to content that falls outside the view area. If essential objects of interest fall in the off-screen region, users need further panning and zooming to access them. As a consequence, spatial tasks that would be fairly easy if all objects of interest were visible at the desired level of detail become difficult and time-consuming, lowering user performance and decreasing satisfaction with mobile applications. This is a significant problem that may nullify

the advantages of anytime-anywhere availability of information, especially in those domains where it is important for the user to rapidly gain situation awareness by glancing at the screen (e.g., in decision-making or coordination activities).

The literature proposes several approaches that can be helpful to mitigate the negative impact of off-screen content during spatial tasks. *Restructuring* an information space into areas of related content that fit as much as possible the display is an idea that researchers have proposed to reduce the amount of panning and zooming needed to view large information spaces on small screens. This is useful in the specific case of web pages, which can be automatically reformatted, e.g. by concatenating all columns, to provide more appropriate viewing modes for mobile devices (Buyukkokten et al., 2000; Chen et al., 2003; Lam and Baudisch, 2005; Roto et al., 2006). However, the approach is unsuitable for information spaces such as maps, whose structure cannot be easily changed without negatively affecting spatial tasks.

Some researchers have focused on reducing the complexity of panning and zooming by implementing *custom pan and zoom mechanisms* that make it easier for users to retrieve relevant content, e.g. by combining scrolling and zooming into a single operation (Robbins et al., 2004; Jones et al., 2005; Burigat et al., 2008a). Although these simplified navigation mechanisms help in

* Corresponding author. Tel.: +39 0432558752; fax: +39 0432558499.

E-mail addresses: stefano.burigat@uniud.it (S. Burigat), luca.chittaro@uniud.it (L. Chittaro).

reducing the effort required in exploring the information space, they actually do not provide any way to make users aware and keep track of off-screen objects during spatial tasks.

Focus + Context visualization is based on displaying an information space at different levels of detail simultaneously, without separating the different views (Leung and Apperley, 1994). Usually, one or multiple focus areas with undistorted content are embedded in surrounding context areas that are distorted to fit into the available screen space. For example, in the Rectangular FishEye View (Rauschenbach et al., 2001), a rectangular focus is surrounded by one or more context belts, appropriately scaled in such a way that less detail is displayed as the distance from the focus increases. The disadvantage of Focus + Context visualizations is that the different scales and distortions they use make it difficult for users to integrate all information into a single mental model and interfere with spatial tasks that require geometric assessments (Baudisch et al., 2002; Nekrasovski et al., 2006).

Unlike Focus + Context, the *Overview + Detail* approach typically displays an overview of an information space and a detail view of a portion of that space simultaneously but in separate views. The overview is usually a small-scale thumbnail of the whole information space and includes a properly positioned graphical highlight (hereinafter, *viewfinder*) of the portion of space that is currently displayed by the detail view. This approach can be potentially useful in supporting spatial tasks since it does not change the structure of an information space and the overview can be used to highlight all objects of interest which are outside the detail view area. However, the feasibility of Overview + Detail visualization on mobile devices has been scarcely investigated and results on its effectiveness, to date, have been conflicting (Büring et al., 2006; Burigat et al., 2008b).

Unlike the other approaches, *Contextual Cues* visualizations have been specifically proposed as a way to provide the user with appropriate information to locate relevant objects even when they are off-screen. In particular, the Contextual Cues approach is based on displaying abstract shapes (or *proxies*) in the border region of the screen to function as visual references to objects of interest that are outside the view area. For example, stylized arrows can be used to point at the location of off-screen objects, highlighting direction information, while size, length, color or other properties of arrows can convey distance information Burigat et al. (2006).

Several Contextual Cues visualizations that differ in the way each proxy conveys direction and distance information to the user have been proposed by researchers in the latest years. However, the effectiveness of each proposal has been studied only through experimental comparisons with some other Contextual Cues visualizations, without taking into consideration the alternative approaches we mentioned above, whose relative merits thus remain unknown. As we pointed out before, at least one of these approaches (Overview + Detail) seems to be suitable to support spatial tasks that involve off-screen objects. One of the goals of this paper is to investigate whether Overview + Detail can be indeed useful in providing information about off-screen objects and how this approach compares to well-known Contextual Cues visualizations. The interesting point is that the two approaches convey their information in very different ways that have implications on the mental effort required to carry out spatial tasks. Indeed, while Contextual Cues visualizations force users to determine the spatial configuration of off-screen objects by examining the properties of on-screen proxies, Overview + Detail directly displays the configuration in the overview. This possible advantage of Overview + Detail might however be nullified by the drawbacks of all mobile implementations of the approach, such as the small size of the overview and the difficulty of relating overview and detail views (Chittaro, 2006). The study we present will help clarify this issue.

Another goal of the paper is to close a gap in the analysis of Contextual Cues visualizations, comparing the effectiveness of Scaled Arrows (Burigat et al., 2006) and Wedge (Gustafson et al., 2008). The two techniques are representative of two different philosophies to provide proxy-based information about off-screen objects to the user: Scaled Arrows separately convey direction and distance information, requiring the user to refer to a legend to precisely interpret the latter; Wedge directly conveys information about the exact location of off-screen objects, relying on user's ability to visually complete partial geometric shapes displayed on the screen. The fact that both techniques were found to provide advantages in certain spatial tasks over Halo (Baudisch and Rosenholtz, 2003), the most known Contextual Cues visualization, makes their direct comparison even more interesting.

The paper is organized as follows. Section 2 surveys related work. Section 3 presents the three techniques compared in our study. Section 4 describes the experimental evaluation and reports results, which are then discussed in Section 5. Finally, Section 6 presents conclusions.

2. Related work

In this section, we present and discuss the literature on the two approaches that are the subject of our study: Overview + Detail and Contextual Cues.

2.1. Overview + Detail

Overview + Detail visualizations display one or multiple overviews of an information space as small-scale thumbnails, together with a detailed view of the specific portion of space highlighted by the viewfinder (Plaisant et al., 1995). Studies of Overview + Detail on desktop computers found that the overview can be an effective tool to support search tasks in an information space (Beard and Walker, 1990; North and Shneiderman, 2000; Pietriga et al., 2007) and can provide benefits to users in terms of information acquisition during navigation (Hornbaek and Frokjaer, 2003). There is also a good amount of evidence in support of user preference for Overview + Detail over other visualizations, even in those studies which found Overview + Detail to be worse than other approaches in terms of performance (Hornbaek et al., 2002). However, the Overview + Detail approach is problematic on mobile devices: fitting overview and detail view on the screen while guaranteeing readability of their content is difficult. Moreover, the screen space that can be assigned to visualize overviews is typically insufficient to allow the user to easily relate them to the detail view (Chittaro, 2006). Very few empirical studies have been carried out to determine how mobile device limitations affect the design and use of Overview + Detail visualizations. Roto et al. (2006) proposed a solution to visualize web pages on small screens by dynamically reformatting pages and overlaying their detail view with an overview of the whole page. The authors found that their approach scored better in usability ratings and user preference compared to a traditional mobile browser that reformatted and displayed web page content in a single column. However, it was impossible to determine whether the results were due to the reformatting technique, the overview, or to the combination of the two factors. Büring et al. (2006) report the results of a user study in which participants performed search tasks on scatterplots by using a detail-only zooming visualization and an Overview + Detail visualization on a PDA. Results revealed that participants with high spatial ability solved tasks significantly faster with the zooming interface. This may suggest that, on small screens, a larger detail view can outweigh the benefits gained from the presence of an overview window. However, results could also have been affected

by the availability of labels on the scatterplot which could have provided users with additional navigation cues beyond those of the overview. In a recent study where we compared two Overview + Detail implementations with a traditional zooming interface on three types of information space (maps, diagrams, and web pages) (Burigat et al., 2008b), we found that an overview brings enough benefit to justify the space used for it on mobile screens if it highlights relevant semantic information that users can exploit during navigation, especially when the considered information space does not include appropriate orientation cues. Thus, the possibility of highlighting objects of interest in the overview might make Overview + Detail useful in supporting spatial tasks involving off-screen objects.

2.2. Contextual cues

Contextual Cues visualizations are explicitly aimed at providing information about the location of objects of interest which are outside the view area. This is obtained by using proxies, i.e. abstract shapes that represent the objects and are overlaid onto the border region of the display window. Zellweger et al. (2003) introduced CityLights in the desktop domain as one of the first examples of this approach. CityLights is a technique that conveys the size of off-screen objects (in the original work, windows in a spatial hypertext system) by projecting them as lines onto the display border. CityLights also offers a coarse representation of object distance by using colors to encode different distance ranges. EdgeRadar (Gustafson and Irani, 2007) is a follow-up to CityLights that reserves a band along the screen border to represent off-screen space and conveys distance information through a mechanism based on the position of small proxies in the band: the closer a proxy is located with respect to the edge of the display, the farther the distance of the corresponding off-screen object. Both CityLights and EdgeRadar use a symbolic representation of distance that requires to provide users with a legend to fully understand the mapping between distance cues and actual distance of off-screen objects. Baudisch and Rosenholtz (2003) introduced a solution called Halo to overcome the need for such a legend. Halo shows the location of off-screen objects by surrounding them with circles that are just large enough to reach into the border region of the display window. Users can thus estimate the off-screen location of objects by looking at the position and curvature of the portion of circles visualized on-screen. An empirical study revealed that Halo enables users to complete map-based route planning tasks faster than a technique based on displaying arrows coupled with labels for distance indication. In a similar study, we compared Halo to Scaled and Stretched Arrows that encode distance as size and

length of arrows, respectively (Burigat et al., 2006). Our results show that Halo improves performance when precise distance is required while Scaled Arrows are more effective than Halo when users need to know the relative distance order of off-screen objects. Moreover, our findings showed that arrow-based visualizations can outperform Halo in the case of cluttered configurations where several off-screen objects must be taken into account. Irani et al. (2006) used oval halos to reduce the issue of overlap and clutter among proxies in Halo but found that the distortion negatively affects distance awareness and prevents users from accurately locating off-screen objects. Gustafson et al. (2008) proposed a technique called Wedge to avoid overlap and clutter in a more effective way. Wedge uses acute isosceles triangles instead of circles to point at off-screen locations. The tip of each triangle coincides with an off-screen object while the base and part of the two legs of the triangle are displayed on screen to convey location information. To remove overlap, triangles are rotated away from each other using an iterative algorithm. The user study in Gustafson et al. (2008) reports that users were significantly more accurate in precisely locating off-screen objects when using Wedge than when using Halo, especially when off-screen objects were clustered into corners.

3. The considered visualizations

Fig. 1 shows examples of the three visualizations we compared in our study: Scaled Arrows, Wedge, and Overview + Detail.

Scaled Arrows, proposed in Burigat et al. (2006), follow the Contextual Cues approach and convey direction and distance information about off-screen objects. Arrow orientation is used to encode direction while arrow size conveys distance (Fig. 2). The size of arrows is inverse-linearly proportional to the object distance from the screen: the larger the arrow, the closer to the screen is the off-screen object. Therefore, the visualization associates those off-screen objects that are closer to the view area to graphical elements that are easier to notice. A drawback of the visualization is that users need a legend to precisely map arrow size into actual distance of off-screen objects. Once users have estimated the distance of an object, they can mentally determine its position through a projection along the direction pointed by the considered arrow. Building an accurate spatial configuration of off-screen objects is therefore complex but Scaled Arrows are effective, as shown in Burigat et al. (2006), when the user needs to qualitatively compare multiple objects, e.g. to order them with respect to their distance from a given on-screen point.

Wedge (Gustafson et al., 2008) is a Contextual Cues visualization that conveys location information about off-screen objects through acute isosceles triangles. On-screen, users see the base



Fig. 1. The three visualizations considered in the study: (a) Scaled Arrows, (b) Wedge, (c) Overview + Detail.



Fig. 2. Scaled Arrows. Each arrow directly conveys direction of the corresponding off-screen object. Arrow size conveys distance: the larger the arrow, the closer to the screen is the off-screen object.

as well as part of the two legs that point towards the off-screen object whose location coincides with the tip of the triangle (Fig. 3). Wedge is based on the theory of amodal completion, which suggests that the human visual system will complete parts of an object even when the object is only partially visible (Elder and Zucker, 1993). More specifically, the design of Wedge is based on the local process of visual completion, which suggests that the visual system completes the occluded part by connecting the extensions of the visible contours (Sekuler et al., 1994). Wedge should thus enable users to accurately determine the location of each off-screen object even if building the spatial configuration of all objects still requires, as in Scaled Arrows, to mentally combine the information provided by all proxies. Another key feature of Wedge is that triangles have three degrees of freedom. It is possible to change rotation, aperture, and intrusion on the display window of each triangle while keeping it pointed at the same location. This is essential in avoiding overlap with other triangles as well as improving location accuracy. In our study, we used the approach defined in Gustafson et al. (2008) to determine aperture and intrusion of each triangle and, when needed, we employed rotation to remove overlaps.

The Overview + Detail visualization we considered displays the overview of the information space as a small thumbnail that covers about 10% of the screen at the bottom right corner of the detail view (Fig. 4). The overview highlights objects of interest as colored dots and contains a viewfinder that shows which is the portion of space displayed in the detail view. Unlike the two Contextual Cues visualizations, Overview + Detail does not require users to mentally build the spatial configuration of objects of interest because the configuration is clearly provided in the overview. However, the small size of the screen might make it difficult to relate the




Fig. 3. Wedge. Each wedge is a triangle whose legs intrude into the screen. Users should triangulate the off-screen object location by mentally completing the partially visible triangles.



Fig. 4. Overview + Detail. The overview displays the spatial configuration of all off-screen objects.

overview and detail view, which might have negative consequences on user's ability to carry out some spatial tasks.

4. User study

The main objective of the study was to compare the effectiveness of Scaled Arrows, Wedge, and Overview + Detail in solving the off-screen objects problem. Overview + Detail has never been studied before as a possible solution to the problem. Scaled Arrows and Wedge were individually compared with Halo in previous studies (Burigat et al., 2006; Gustafson et al., 2008), but no study ever compared arrows with Wedge. Moreover, we were interested in studying the possible effects of a moderate increase in the number of off-screen objects on each visualization. To obtain results that could be compared to our previous study (Burigat et al., 2006) which took the number of objects into consideration, we required users to carry out spatial tasks with configurations of 5 and 8 off-screen objects (see Fig. 5).

Considering the specific features of the three visualizations, our hypotheses in the study were the following:

- Overview + Detail enables users to be faster than both Contextual Cues visualizations in carrying out spatial tasks that require knowledge of the spatial configuration of off-screen objects. This hypothesis is supported by the fact that Overview + Detail provides direct visual access, albeit on a small-scale, to the configuration of off-screen objects, while both Contextual Cues visualizations require users to mentally build the configuration of off-screen objects through examination of the graphical features of proxies.
- Overview + Detail and Wedge allow users to be more accurate than Scaled Arrows in determining the location of off-screen objects. This hypothesis is supported by the fact that Scaled Arrows provide only a symbolic representation of distance while both Overview + Detail and Wedge enable users to accurately find out where off-screen objects are located, by directly showing the position on the overview in the first case and through amodal completion in the second case.
- Regardless of the visualization, users should be slower in carrying out tasks as the number of off-screen objects increases. There should also be a negative effect of the number of objects on user accuracy in tasks requiring comparisons of multiple off-screen objects. However, there should be no effect on user accuracy in determining the precise location of off-screen objects since each object can be examined independently of the others.

4.1. Participants

Twenty-four subjects (9 female, 15 male) were recruited by direct contact to participate in the study. Sixteen of them were

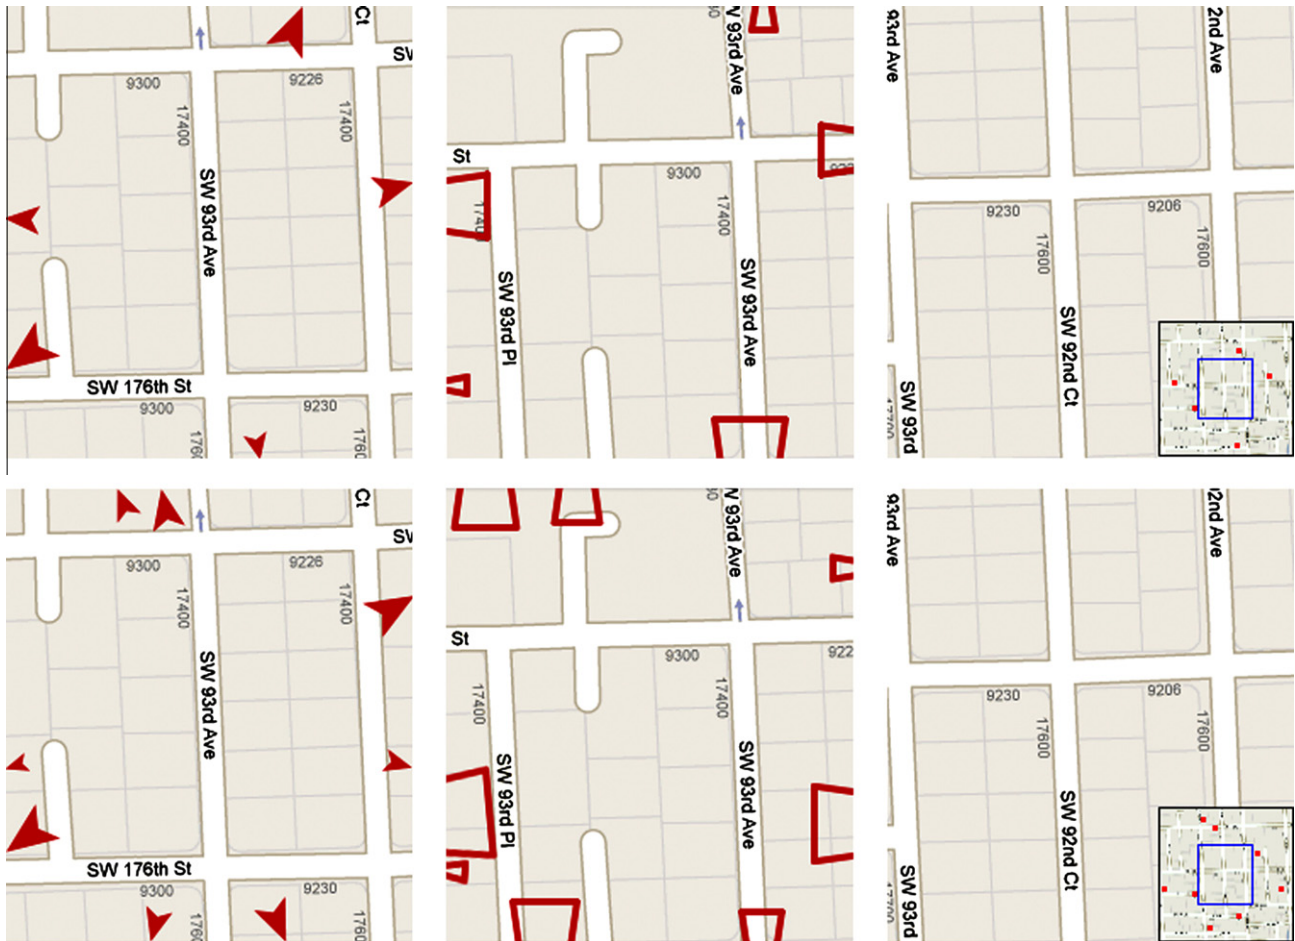


Fig. 5. Examples of configurations containing 5 and 8 off-screen objects, represented with each of the three considered visualizations.

undergraduate students from various backgrounds (7 Computer Science, 5 Architecture, 2 Humanities, 2 Engineering) and the remaining 8 were people from diverse occupations (5 self-employed persons, 3 factory workers). Their age ranged from 20 to 59, averaging at 29. All users had high familiarity with mobile devices, most users had at least some familiarity with the use of geographic maps in general (8 high familiarity, 10 medium, 5 low, 1 no familiarity), most users had low or no familiarity with the use of geographic maps on mobile devices (4 high familiarity, 4 medium, 4 low, 12 had never used mobile maps).

4.2. Materials

The study was carried out on an Asus P535 Windows Mobile phone featuring a 520 MHz processor and a 2.8-in. touchscreen with a resolution of 240×320 pixels. During the evaluation, a 240×268 area in the middle of the screen displayed the considered visualization while two menu bars were displayed at the top and at the bottom of the screen. Fig. 6 shows an example of the Scaled Arrows visualization on the phone. All maps we used in the study were structurally similar and represented city areas (unknown to users) with a number of intersecting streets. Participants interacted with the phone using the stylus.

4.3. Tasks

We considered a mobile scenario in which the user had to carry out spatial tasks involving objects of interest on a city map. We set the zoom level so that only a limited portion of the map could be

displayed at once on the screen and the scenario thus involved objects of interest that fell outside the currently displayed area. In this way, the user had to rely on off-screen location visualization to quickly complete tasks. In particular, we considered the following tasks:

- *Closest*: point out the off-screen object that is closest to the screen border. The screen displays the area the user is currently in (considering the center of the screen as user's position) but all relevant objects are off-screen. This is one of the most common spatial tasks for map users (e.g., tourists who need to reach a nearby point of interest such as a monument or a restaurant) and has been considered in all past evaluations of off-screen location visualizations. In this task, users were required to provide their answers by tapping on (or very near) the graphical element associated to the closest off-screen object.
- *Order*: order all off-screen objects in increasing distance from the screen border. This is a relatively complex spatial task that requires users to compare the distance of all off-screen objects. To carry out the task, users had to tap in distance order on (or very near) all graphical elements associated to off-screen objects.
- *Cluster*: point out the pair of off-screen objects which are closest to each other. This task is interesting because of its spatial complexity, due to the need for users to reason in terms of the possible pairs of off-screen locations. To complete the task, users had to tap on (or very near) the two graphical elements associated to the selected pair of off-screen objects of interest.




Fig. 6. The phone we used in the study.

- *Locate*: mark the off-screen location of each off-screen object on a printed version of the visualization. In previous studies about off-screen location visualization, this is the task that best revealed the effectiveness, in terms of accurate location estimation, of visualizations based on amodal completion (i.e., Halo and Wedges). Users had to carry out the task on a sheet of paper containing a printout of the visualization to analyze. The visualization was centered so that the area where the user had to mark the location of each off-screen object was left blank. During the task, the visualization was displayed on the mobile phone as well.

The tasks we employed are very similar to those used in the related literature (Burigat et al., 2006; Gustafson et al., 2008) and differ among them in the amount of information about off-screen objects needed to carry them out. Two tasks (Closest and Order) rely on the capability of comparing the distance of objects of interest from the screen border, one task (Locate) requires users to accurately estimate the location of each off-screen object, and one task (Cluster) requires knowledge of the exact spatial configuration of all objects. Since there are many different tasks which may involve off-screen objects, a focus on the knowledge they require rather than on the specific task may help in generalizing the results while keeping studies manageable.

In the first three tasks, which were carried out directly on the mobile phone, a sound was played every time users tapped on the screen and a small circular glyph was displayed in the tapped position (Fig. 7). These aids were aimed at providing feedback to users, helping them understand when and where they had actually tapped the screen. This was useful to avoid situations in which the




Fig. 7. During task execution, small circular glyphs provided users with visual feedback about tapped points.

system did not log a tap because the user tapped too lightly on the screen or situations in which the user was unsure if she had already selected a target or not.

4.4. Experimental design and procedure

The study was based on a 3×2 within-subjects factorial design with two factors, *Visualization* (Scaled Arrows, Wedge, or Overview + Detail) and *Number of objects* (5 or 8 off-screen objects). Participants were individually briefed about the nature of the experiment and were asked to fill in a short demographic questionnaire which contained also questions about the degree of familiarity with mobile devices and geographic maps (paper and digital ones). Then, the experimenter described the three visualizations and the tasks to be performed, carefully explaining the mapping between proxies and position of the corresponding off-screen objects. It was stressed that Scaled Arrows associated larger proxies to off-screen objects that were closer to the view area while Wedge did the opposite. After this introductory phase, users carried out 24 experimental tasks (3 visualizations \times 2 sets of objects \times 4 types of task), each one preceded by an appropriate training task to let users familiarize with the considered combination of conditions (but using a different configuration of off-screen objects). During training, users could talk with the experimenter to clarify possible doubts and they were also informed whether they were using an incorrect mapping between proxies and off-screen objects. To start Closest, Order, and Cluster tasks, users were required to tap on a “Start Task” button that was initially displayed on the screen. Each task ended when users tapped on the last target. Participants were also required to rank visualizations according to their preference after they had concluded all tasks of a specific type (Closest, Order, Cluster, or Locate), separately for each of the two levels of the “Number of objects” factor. The average duration of the test, which was carried out in a lab setting, was approximately 30 min.

To avoid any sequence or learning effects, the order of task, Visualization and Number of objects were counterbalanced. The

spatial configuration of off-screen objects was also systematically varied and controlled across the different conditions. More specifically, we manually prepared 8 master configurations (4 for each level of the Number of objects factor) in such a way that there was only one correct answer to each task (e.g., there was only one closest object in the Closest task). Objects were distributed at different distances on all four sides of the view area. In any given task, we ensured that configurations had exactly the same complexity with all three visualizations by using one of the master configurations with the first visualization, a mirrored version of the configuration with the second visualization and a flipped version of the configuration with the third visualization. We also ensured that there was no fixed association between task and master configuration. Due to the large visual difference among the visualizations, no user noticed the use of mirrored and flipped versions of configurations, as we assessed after the user had completed all tasks.

For each user, we automatically logged task completion time and tapped points for the Closest, Order, and Cluster tasks. Data analysis was carried out on task completion times and error rates derived from tapped points (for the Closest, Order and Cluster tasks). For the Locate task, we analyzed accuracy in the Locate task (derived by distance measurements on the paper sheets administered to users).

4.5. Results

4.5.1. Task completion time

Figs. 8–10 show mean completion times for the Closest, Cluster, and Order tasks, for all six possible combinations of the two within-subjects factors (visualization, number of objects). The three levels of the visualization factor are abbreviated as ScA (for Scaled Arrows), Wed (for Wedge), and OD (for Overview + Detail). For the Order task, since ordering 8 objects necessarily requires more time than ordering 5 objects, dividing the total time by the number of objects provides more meaningful data for the analysis. Fig. 9 thus provides mean completion times divided by the number of objects (5 or 8).

Task completion times were subjected to the Shapiro–Wilk test of normality prior to further analysis. The test revealed moderate deviations from the normal distribution for all three tasks and data was normalized using a log transformation. A two-way repeated measures analysis of variance (ANOVA) was then employed on the log-transformed times.

For the Closest task (Fig. 8), ANOVA did not reveal a significant interaction between visualization and number of objects ($F(2,46) = 0.16, p = 0.85$). While users took more time to complete the task with 8 off-screen objects than they did with 5 objects with all visualizations, the main effect of number of objects did




Fig. 8. Mean completion times (with standard error bars) for the Closest task.




Fig. 9. Mean completion times (with standard error bars) for the Order task, divided by the number of objects.




Fig. 10. Mean completion times (with standard error bars) for the Cluster task.




Fig. 11. Mean error rate (with standard error bars) in the Closest task.

not reach significance ($F(1,23) = 4.04, p = 0.056$). No significant main effect of visualization was detected either ($F(2,46) = 0.14, p = 0.87$).

For the Order task (Fig. 9), the ANOVA did not reveal a significant interaction between visualization and number of objects ($F(2,46) = 0.73, p = 0.47$). No main effect of number of objects was detected ($F(1,23) = 0.12, p = 0.73$) while there was a main effect of visualization ($F(2,46) = 7.59, p < 0.005$). Since there was no interaction, the main effect of visualization was further investigated using Tukey’s post-hoc test on marginal means. The analysis revealed that users were significantly slower in completing the task with Overview + Detail than they were with Wedge ($q = 5.54, p < 0.05$) and with Scaled Arrows ($q = 3.52, p < 0.05$).

For the Cluster task (Fig. 10), a significant interaction was detected ($F(2,46) = 3.55, p < 0.05$). No significant main effect was found for number of objects ($F(1,23) = 3.40, p = 0.078$) while a significant main effect was found for visualization ($F(2,46) = 14.11, p < 0.001$). To investigate the interaction, we compared cell means using Tukey's test to look for the effects of visualization at each level of number of objects and t -tests to look for the effects of number of objects at each level of visualization. With 5 off-screen objects, the analysis revealed a statistically significant difference between Wedge and Overview + Detail ($q = 5.05, p < 0.05$). With 8 off-screen objects, Tukey's test revealed a statistically significant difference between Wedge and Overview + Detail ($q = 4.51, p < 0.05$) and between Scaled Arrows and Overview + Detail ($q = 6.50, p < 0.05$). These results also show that the main effect of visualization was not consistent across all levels of number of objects. Users were significantly faster in completing the Cluster task with Overview + Detail than they were with Wedge, regardless of the number of off-screen objects, but they were significantly faster with Overview + Detail than with Scaled Arrows only with 8 off-screen objects. The t -tests revealed that users were significantly faster in completing the Cluster task with 5 off-screen objects than they were with 8 off-screen objects when using Scaled Arrows ($t = 3.54, df = 23, p < 0.005$) while there were no statistically significant differences when using Wedge or Overview + Detail.

4.5.2. Error

Figs. 11–13 show error rates for the Closest, Order and Cluster tasks. Error rates indicate the percentage of users who gave a wrong answer, i.e. did not correctly locate the closest off-screen object, the correct order of off-screen objects, or the pair of off-




Fig. 12. Mean error rate (with standard error bars) in the Order task.




Fig. 13. Mean error rate (with standard error bars) in the Cluster task.




Fig. 14. Mean distance error (with standard error bars) in the Locate task, divided by the number of objects.

screen objects closest to each other, respectively in the Closest, Order and Cluster tasks.

As with task completion times, the Shapiro–Wilk test of normality we performed prior to further analysis revealed deviations from the normal distribution. Since no transformation could normalize the data, we employed a non-parametric procedure for mixed models, the ANOVA-Type Statistic (ATS) (Brunner and Munzel, 1999), to analyze main and interaction effects.

For the Closest task (Fig. 11), the ATS revealed no significant main effect of number of objects ($ATS = 0.81, p = 0.37$) and no significant main effect of visualization ($ATS = 2.09, p = 0.15$). A plot of the cell means revealed a total lack of interaction between the two factors, highlighted by perfectly parallel lines (in such condition, the ATS value is close to 0 but the exact p -value cannot be computed).

For the Order task (Fig. 12), the ATS revealed no interaction effect ($ATS = 2.3, p = 0.11$), a significant main effect of number of objects ($ATS = 31.04, p < 0.001$), with a much higher error rate for configurations of 8 objects, and no significant main effect of visualization ($ATS = 2.93, p = 0.054$).

For the Cluster task (Fig. 13), the ATS revealed no interaction effect ($ATS = 0.36, p = 0.69$), no main effect of number of objects ($ATS = 0.096, p = 0.76$), but a significant main effect of visualization ($ATS = 3.75, p < 0.05$). We used Dunn's post-hoc test to investigate the main effect of visualization. The analysis pointed out that the error rate with Overview + Detail was significantly lower than both the error rate with Scaled Arrows ($p < 0.05$) and the error rate with Wedge ($p < 0.05$).

Fig. 14 shows the mean distance error in the Locate task. Distance error was measured as the Euclidean distance, in pixels, between the subject's location estimate and the actual location of an off-screen object. As for times in the Order task, we divided the total error by the number of objects to obtain more meaningful data for the analysis. We used a square root transformation on the data to correct the moderate deviation from the normal distribution revealed by the Shapiro–Wilk test of normality. A two-way repeated measures ANOVA was then employed on the transformed data. The analysis pointed out a significant interaction effect ($F(2,46) = 5.26, p < 0.01$), as well as a significant main effect of number of objects ($F(1,23) = 8.65, p < 0.01$) and visualization ($F(2,46) = 6.76, p < 0.005$). To investigate the interaction, we compared cell means using Tukey's test to look for the effects of visualization at each level of number of objects and t -tests to look for the effects of number of objects at each level of visualization. Tukey's test revealed that users were significantly less accurate with Scaled Arrows than they were with Wedge ($q = 4.17, p < 0.05$) and Overview + Detail ($q = 6.30, p < 0.05$) with 8 off-screen objects. Since no such effect was found with 5 off-screen objects, the main effect

of visualization was not consistent across all levels of the number of objects variable. The t -tests revealed that users were significantly more accurate in the Locate task with 5 off-screen objects than they were with 8 off-screen objects when using Scaled Arrows ($t = 2.86$, $df = 23$, $p < 0.01$) and Wedge ($t = 2.69$, $df = 23$, $p < 0.05$). However, the main effect of number of objects was not consistent across all levels of visualization since no statistically significant difference in accuracy was found for Overview + Detail between 5 and 8 off-screen objects conditions.

4.5.3. Subjective preference

Figs. 15–18 show subjective preference for the three visualizations in the Closest, Order, Cluster, and Locate tasks. To analyze the data, we employed the non-parametric ATS, followed by Dunn's test where appropriate. Since users were asked to rate the three visualizations from the best to the worst, we assigned a score of 3, 2, 1 respectively to the first, second, and third visualization. An appropriate fractionary score was assigned to draws, which were allowed.

The analysis revealed a significant main effect of visualization in the Cluster task ($ATS = 23.2$, $p < 0.001$). Dunn's post-hoc test pointed out a statistically significant difference in preference between Overview + Detail and Scaled Arrows ($p < 0.05$) as well as between Overview + Detail and Wedge ($p < 0.05$), with users preferring the first visualization in both cases. No significant main or interaction effect emerged in the other tasks.

5. Discussion

Overall, results of the analysis were multifaceted and only partially met our initial hypotheses.




Fig. 15. Mean preference (with standard error bars) in the Closest task.




Fig. 16. Mean preference (with standard error bars) in the Order task.




Fig. 17. Mean preference (with standard error bars) in the Cluster task.




Fig. 18. Mean preference (with standard error bars) in the Locate task.

In the Closest task, no visualization had a significant advantage over the others. While this is not a surprising result since the task was simple and did not require users to be spatially aware of the configuration of off-screen objects, it is interesting because the different strategies users had to employ to carry out the task (scan the visualization to identify the smallest proxy in Wedge, scan the visualization to identify the biggest proxy in Scaled Arrows, look at the overview to identify the dot nearest to the viewfinder in Overview + Detail) did not have an impact on the outcome. The ease of use of the three visualizations in this task is also confirmed by the very low error rate in all conditions. All visualizations were found to be quite effective for the Closest task also in case of a moderate increase of the number of off-screen objects. However, it must be noted that users consistently took more time to carry out the task in all conditions when the number of off-screen objects was higher. It is also interesting to note that, in this task, some users selected the farthest off-screen object rather than the closest (3 with Scaled Arrows and 2 with Wedge) despite the training. This could be due to the difficulty in changing the strategy to carry out the task (finding the largest proxy in Scaled Arrows and finding instead the smallest proxy in Wedge) when switching between the two techniques. This kind of event did not occur with Overview + Detail.

Contrary to our first hypothesis, in the Order task users were significantly slower with Overview + Detail than they were with Wedge and Scaled Arrows. This task can be considered as a more complex version of Closest, requiring users to identify the closest off-screen object, then the closest among the remaining off-screen objects and so on up to the farthest off-screen object. As with Closest, users do not need to be spatially aware of the configuration of off-screen objects. Thus, a possible explanation of the result is that it is easier for users to directly compare the size of proxies with

Wedge or Scaled Arrows than it is to compare the distances of dots from the viewfinder in a small-scale overview. This is an example of the drawbacks of small overviews on mobile devices, which can nullify the advantage of having direct visual access to object configurations. We also found no significant effect of number of objects on task completion time (divided by the number of objects). While this outcome was in line with the corresponding one in the Closest task in terms of statistical significance, it is remarkable that there was practically no difference in the mean times between 5-objects and 8-objects conditions with Scaled Arrows and Overview + Detail, and a very small difference with Wedge. This is surprising, since we expected an increase in the time needed to order each off-screen object with 8 objects similar, in percentage, to the one we found in the Closest task. Moreover, the absolute time to carry out the Closest task was much higher than the absolute time to order a single off-screen object in the Order task. This might mean that, once a user has identified the closest off-screen object, it is much easier to find subsequent off-screen objects in order of distance. Finally, we did find a significant effect of number of objects on error rate, which greatly increased as the number of off-screen objects increased. Probably, cluttered configurations make it more likely to have off-screen objects at similar distance from the display window, thus raising the difficulty of the task. It is also interesting to note that, while Halo was found to be less effective than Scaled Arrows in one of our previous studies (Burigat et al., 2006), Wedge did instead outperform Scaled Arrows in the Order task of this study, which indirectly confirms how Wedge succeeded in improving Halo.

The Cluster task revealed the effectiveness of Overview + Detail when it is important to know the spatial configuration of off-screen objects. In this task, which is probably the most complex of the four since it requires users to reason in terms of the location of pairs of off-screen objects, users were significantly faster and were more accurate with Overview + Detail than they were with Wedge, regardless of the number of off-screen objects. The comparison with Scaled Arrows was likewise remarkable. The most likely explanation for these results is that with Overview + Detail users do not need to build an internal mental image of the configuration of off-screen objects before comparing the distance between pairs of objects because the configuration is externally visible in the overview. This mental operation is instead required for Scaled Arrows and Wedge. We did not find an effect of number of objects on task completion times with Wedge and Overview + Detail but users were significantly slower with 8-objects configurations than they were with 5-objects configurations with Scaled Arrows, probably because it is much more difficult to build a mental image of the configuration when a visualization provides only qualitative distance information.

In the Locate task, we expected Overview + Detail and Wedge to be more effective than Scaled Arrows in terms of user accuracy, regardless of the number of off-screen objects, because they provide more powerful means to accurately locate off-screen objects (a visual representation of the configuration of objects in the first case, amodal completion in the second case) while Scaled Arrows provides only a qualitative representation of distance. We also expected the number of objects to have no effect on the results because users could focus on each object independently of the others to find out its location and were given no time pressure to finish the task. However, our two expectations were only partially confirmed. Overview + Detail and Wedge were more effective than Scaled Arrows in the case of 8-objects configurations only. The lack of difference among the three visualizations for 5-objects configurations is surprising and difficult to explain. Unfortunately, we did not measure the time needed by users to complete the task (since it would have been difficult to accurately control start and end time in this paper-based task) and thus cannot check if increased

accuracy came at the expense of task completion time. The other surprising result was the negative effect of number of objects on accuracy with Scaled Arrows and, in particular, Wedge. A possible explanation for this result is that as users made marks on paper, they employed those marks to guide the identification of next locations, thus possibly compounding the error. While using previous marks could be a reasonable strategy for Scaled Arrows, it is not for Wedge, which allows to use each proxy independently of the others to find out the off-screen object location. We could also consider that with 8 off-screen objects proximity among proxies increases (see Fig. 5) and this could have negatively affected the process of projection (with Scaled Arrows) and closure (with Wedge). Further investigation is needed to more precisely determine the causes of these results.

Finally, subjective preference was in line with performance results in the Cluster task while no statistically significant differences emerged in the Order and Locate tasks. Strangely, users highly rated Overview + Detail in the Order task despite their low performance with the technique (compared to Wedge) both in terms of completion time and accuracy. This may be due to the fact that users prefer having direct visual access to the configuration of off-screen objects even if the small size of the overview makes it actually difficult to easily extract accurate information.

6. Conclusions

Overall, results of the study show that there is no single best solution to support users in carrying out different spatial tasks on mobile devices when relevant objects are off-screen. In particular, we found that Overview + Detail on mobile devices is a useful solution for the off-screen objects problem and is more effective than Contextual Cues visualizations when the user needs to reason in terms of the spatial configuration of off-screen objects (as in the Cluster task). However, Wedge is more effective when only distance information of all off-screen objects is important (as in the Order task). The two visualizations seem instead to be equally good solutions when the location of individual off-screen objects must be accurately estimated or when the user is simply interested in the closest object. In this latter case, which is one of the most common in traditional mobile map applications, Scaled Arrows are also appropriate. Designers of mobile applications that support activities in which the user needs to gain spatial awareness of the information space (e.g., decision support systems, geographic information systems) have instead no simple choice. As we saw in the results, choosing the wrong visualization for a certain task impacts user performance. If we focus only on the error metric, which for some applications is more important than the time metric, and also consider that time differences in the order of a few seconds may be acceptable, we find that there is a 25–30% performance difference between Wedge and Overview + Detail in the Order (where Wedge is better) and Cluster (where Overview + Detail is better) tasks. Thus, if both types of task need to be supported, designers have to provide both visualizations or be ready to pay a performance penalty. In such cases, other criteria may play a role in the choice. For example, if orientation support is needed, then Overview + Detail might be a better solution (Burigat et al., 2008b). Finally, we found that even a small increase in the number of off-screen objects negatively affects user performance, in particular error rate, and that Scaled Arrows seems to be the visualization that suffers the most in more cluttered conditions. Further studies are needed to look into the effect of higher numbers of off-screen objects. However, with a very large number of off-screen objects, it is likely that all visualizations will fail to provide useful information, requiring refined or alternative solutions to the problem.

Acknowledgment

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

References

- Baudisch, P., Rosenholtz, R., 2003. Halo: a technique for visualizing off-screen locations. In: Proc. Conference on Human Factors in Computing Systems (CHI 2003). ACM Press, pp. 481–488.
- Baudisch, P., Good, N., Bellotti, V., Schraedley, P., 2002. Keeping things in context: a comparative evaluation of focus plus context screens, overviews, and zooming. In: Proc. Conference on Human Factors in Computing Systems (CHI 2002). ACM Press, pp. 259–266.
- Beard, D., Walker, J., 1990. Navigational techniques to improve the display of large two-dimensional spaces. *Behaviour and Information Technology* 9, 451–466.
- Brunner, E., Munzel, U., 1999. Rank-score tests in factorial designs with repeated measures. *Journal of Multivariate Analysis* 70, 286–317.
- Burigat, S., Chittaro, L., Gabrielli, S., 2006. Visualizing locations of off-screen objects on mobile devices: a comparative evaluation of three approaches. In: Proc. Conference on Human–Computer Interaction with Mobile Devices and Services (Mobile HCI 2006). ACM Press, pp. 239–246.
- Burigat, S., Chittaro, L., Gabrielli, S., 2008a. Navigation techniques for small-screen devices: an evaluation on maps and web pages. *International Journal of Human–Computer Studies* 66, 78–97.
- Burigat, S., Chittaro, L., Parlato, E., 2008b. Map, diagram, and web page navigation on mobile devices: the effectiveness of zoomable user interfaces with overviews. In: Proc. Conference on Human–Computer Interaction with Mobile Devices and Services (Mobile HCI 2008). ACM Press, pp. 147–156.
- Büring, T., Gerken, J., Reiterer, H., 2006. 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). ACM Press, pp. 233–240.
- Buyukkokten, O., Garcia-Molina, H., Paepcke, A., Winograd, T., 2000. Power browser: efficient web browsing for PDAs. In: Proc. Conference on Human Factors in Computing Systems (CHI 2000). ACM Press, pp. 430–437.
- Chen, Y., Ma, W., Zhang, H., 2003. Detecting web page structure for adaptive viewing on small form factor devices. In: Proc. World Wide Web Conference (WWW 2003). ACM Press, pp. 225–233.
- Chittaro, L., 2006. Visualizing information on mobile devices. *IEEE Computer* 39, 40–45.
- Elder, J., Zucker, S., 1993. The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research* 33, 981–991.
- Gustafson, S., Irani, P., 2007. Comparing visualizations for tracking off-screen moving targets. In: Proc. Conference on Human factors in Computing Systems (CHI '07). ACM Press, pp. 2399–2404.
- Gustafson, S., Baudisch, P., Gutwin, C., Irani, P., 2008. Wedge: clutter-free visualization of off-screen locations. In: Proc. Conference on Human factors in Computing Systems (CHI '08). ACM Press, pp. 787–796.
- Hornbaek, K., Frokjaer, E., 2003. Reading patterns and usability in visualizations of electronic documents. *ACM Transactions on Computer–Human Interaction* 10, 119–149.
- Hornbaek, K., Bederson, B., Plaisant, C., 2002. Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM Transactions on Computer–Human Interaction* 9, 362–389.
- Irani, P., Gutwin, C., Yang, X., 2006. Improving selection of off-screen targets with hopping. In: Proc. Conference on Human Factors in Computing Systems (CHI 2006). ACM Press, pp. 299–308.
- Jones, S., Jones, M., Marsden, G., Patel, D., Cockburn, A., 2005. An evaluation of integrated zooming and scrolling on small screens. *International Journal of Human–Computer Studies* 63, 271–303.
- Lam, H., Baudisch, P., 2005. Summary thumbnails: readable overviews for small screen web browsers. In: Proc. Conference on Human Factors in Computing Systems (CHI 2005). ACM Press, pp. 681–690.
- Leung, Y.K., Apperley, M.D., 1994. A review and taxonomy of distortion-oriented presentation techniques. *ACM Transactions on Computer–Human Interaction* 1, 126–160.
- Nekrasovski, D., Bodnar, A., McGrenere, J., Guimbretiere, F., Munzner, T., 2006. 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). ACM Press, pp. 11–20.
- North, C., Shneiderman, B., 2000. Snap-together visualization: can users construct and operate coordinated visualizations? *International Journal of Human–Computer Studies* 53, 715–739.
- Pietriga, E., Appert, C., Beaudouin-Lafon, M., 2007. Pointing and beyond: an operationalization and preliminary evaluation of multi-scale searching. In: Proc. Conference on Human Factors in Computing Systems (CHI 2007). ACM Press, pp. 1215–1224.
- Plaisant, C., Carr, D., Shneiderman, B., 1995. Image-browser taxonomy and guidelines for designers. *IEEE Software* 12, 21–32.
- Rauschenbach, U., Jeschke, S., Schumann, H., 2001. General rectangular fisheye views for 2D graphics. *Computers and Graphics* 25, 609–617.
- Robbins, D., Cutrell, E., Sarin, R., Horvitz, E., 2004. ZoneZoom: map navigation for smartphones with recursive view segmentation. In: Proc. International Conference on Advanced visual interfaces (AVI 2004). ACM Press, pp. 231–234.
- Roto, V., Popescu, A., Koivisto, A., Vartiainen, E., 2006. Minimap: a web page visualization method for mobile phones. In: Proc. Conference on Human Factors in Computing Systems (CHI 2006). ACM Press, pp. 35–44.
- Sekuler, A., Palmer, S., Flynn, C., 1994. Local and global processes in visual completion. *Psychological Science* 5, 260–267.
- Zellweger, P., Mackinlay, J., Good, L., Stefik, M., Baudisch, P., 2003. City lights: contextual views in minimal space. In: Proc. Conference on Human Factors in Computing Systems (CHI 2003). ACM Press, pp. 838–839.