The Effects of Artificial Landmarks on Learning and Performance in Spatial-Memory Interfaces

Md. Sami Uddin[1], Carl Gutwin[1], and Andy Cockburn[2]
[1]Computer Science, University of Saskatchewan
Saskatoon, Canada
[2]Computer Science, University of Canterbury
Christchurch, New Zealand
sami.uddin@usask.ca, gutwin@cs.usask.ca, andy@cosc.canterbury.ac.nz

ABSTRACT
Spatial memory is a powerful way for users to become expert with an interface, because remembering item locations means that users do not have to carry out slow visual search. Spatial learning in the real world benefits greatly from landmarks in the environment, but user interfaces often provide very few visual landmarks. In this paper we explore the use of *artificial* landmarks as a way to improve people's spatial memory in spatially-stable grid menus called CommandMaps. We carried out three studies to test the effects of three types of artificial landmarks (standard gridlines, simple anchor marks, and a transparent image) on spatial learning. We found that for small grid menus, the artificial landmarks had little impact on performance, whereas for medium and large grids, the simple anchor marks significantly improved performance. The simple visual anchors were faster and less error-prone than the visually richer transparent image. Our studies show that artificial landmarks can be a valuable addition to spatial interfaces.

Author Keywords
Command selection; expertise; spatial memory; landmarks.

ACM Classification Keywords
H.5.2. Information interfaces (e.g., HCI): User Interfaces.

INTRODUCTION
Spatial consistency is a powerful means for enabling expert performance with user interfaces. By providing stable spatial locations, users can anticipate the location of items and quickly acquire them. Touch-typing is a good example – people can quickly access specific letters without thinking about key locations, and can even form chains of anticipated motor actions that are executed semi-autonomously (e.g., typing the characters of a word while composing the subsequent sentence). When interfaces fail to support spatial consistency, as they often do, users instead need to resort to comparatively slow visual search to find items, negating opportunities for anticipatory action.

CHI 2017, May 06-11, 2017, Denver, CO, USA
© 2017 ACM. ISBN 978-1-4503-4655-9/17/05…$15.00
DOI: http://dx.doi.org/10.1145/3025453.3025497

Many research and commercial interfaces have been explicitly designed to exploit the efficiencies offered by learned, stable, spatial locations. For example, Marking Menus [28] allow large command vocabularies to be quickly accessed through a fluid series of directional gestures. CommandMaps [42] also allow access to large command vocabularies, but they do so by flattening the traditional command hierarchy, assigning each command to a unique spatial location in the display. Third, gestural ShapeWriting [53] allows users to input text on mobile devices by sweeping out an approximate gesture over a series of spatially stable characters on a virtual keyboard.

Although spatially stable interfaces can enable high input efficiency once the user is an expert, the attainment of expertise depends on the user learning, remembering, and efficiently recalling item locations. Learning to touch type, for instance, typically consumes months of training, and the skill is refined for years. Few office workers, however, would be willing to engage in such dedicated training to become proficient with a new user interface.

There are therefore important research questions in determining effective methods to assist users in learning and recalling spatial locations. For example, previous studies have examined the learning benefits derived from promoting 'deep encodings' [14] – by removing the continual availability of visual feedback, users are forced to actively engage their spatial memory (rather than rely on visual search) which has been shown to improve users' recollection of the location of abstract icons [17], the location of keys on a new keyboard layout [13], and the shape of command gestures [4,7]. However, these approaches intentionally make interaction harder for novice and intermediate users, which may be acceptable for those with a desire to become expert, but will frustrate many others.

An alternative approach for facilitating spatial learning is motivated by a real world mechanism that novices and experts use to augment spatial memory – landmarks. Landmarks are readily identifiable features in space that are easily discriminated from their surrounds [31] and which serve as an orientation point for spatial actions in a familiar or unfamiliar environment. The use of landmarks has been frequently examined in first-person navigation through 3D virtual environments, such as virtual and augmented reality [16,40]. However, there has been comparatively little research into how landmarking features can facilitate object spatial memory in the static and substantially 2D layouts that dominate mobile and desktop interfaces.

In striving for uncluttered and visually appealing user interfaces, contemporary designs often contain few graphical embellishments. While this may improve aesthetics, it also creates a void of potential landmarks that users might otherwise have employed to assist the formation and use of spatial memory. Other interfaces, in contrast, are heavily populated with features that could be used as landmarks. For example, desktop wallpaper images provide a backdrop that may help users memorize the location of icons. Opportunities for leveraging visual embellishments to assist interaction have been examined in previous work – for example, in observing the disparity between clean user interfaces and messy, dog-eared paper documents, Hill et al. [24] proposed the use of *edit-wear* and *read-wear* to graphically augment interface elements with traces of the user's activity. Hill suggested that these augmentations could indicate frequent and recent activity – but few studies have examined the role that visual information and landmarks can play in assisting spatial interaction with user interfaces.

The three studies reported in this paper examine the influence that different forms of artificially-added landmarks have on spatial learning and recall. Experimental tasks involved retrieving items from a grid-menu of alternatives, similar to CommandMaps [42]. The number of candidate items increased across the three studies (8x8, 8x12, and 10x16). Each study compared item location learning across blocks using three forms of landmarking assistance: a *standard* unadorned grid of icons that used only gridlines for background landmarking; a grid augmented with visual *anchors* in the form of gray backgrounds for a few items in the grid (the intention being to provide clear landmarked reference points for the user); and an *image* background that used a transparent overlay image (of the Taj Mahal) plus background gridlines. The abstract landmarks of the *anchored* condition provide highly distinct spatial demarcation, and the *image* condition offers semantically meaningful features (for example, 'the icon by the right turret') – both of these landmarking aids could potentially offer spatial memory advantages. Our overall hypothesis, then, is that the anchor points and the overlay image will assist learning the locations of items in the grid.

Results showed that *anchor* landmarks were most effective in improving users' spatial memory and performance. There was no difference between the techniques in small grids, but with larger grids, error rates and subjective preferences all favored the *anchor* condition.

The studies provide three main contributions. First, we show that for smaller spatial interfaces, artificial landmarks over and above a basic grid offer little benefit. Second, we demonstrate that as interfaces grow larger, the value of artificial landmarks increases significantly. Third, we provide empirical evidence about spatial learning and spatial retrieval that can assist designers as they build future interfaces based on spatial memory.

RELATED WORK

Interfaces for Improved Selection Performance

From entertainment applications to office work, command selection is one of the fundamental tasks that users perform. Selection performance in these interfaces is dependent on two operations. First, users must find a specific command among those available, and second, they must execute that command by pointing to it with a pointing device. Generally, pointing time depends on target width and distance (i.e., Fitts' Law). The time to find a command, however, is related to users' familiarity with the interface [12]. Inexperienced users must rely on slow visual search, but knowledgeable users can skip this step [23] and simply recall the command's location – speeding up performance.

Considerable research has examined methods to improve performance in both of these stages. Alternative command organizations are one main approach: for example, to reduce pointing time, *pie menus* [10] place the commands in a circle around the cursor upon invocation. *Marking menus* [26] use a similar radial organization, but also allow experts to perform pre-emptive gestural selections. Other approaches attempt to flatten command hierarchies to reduce the fixed costs of navigating between levels of the hierarchy (e.g., CommandMaps [42] and FastTap [21] use grid approaches, and other techniques orient items around a user's hand [49]). Keyboard-based shortcuts (i.e., hotkeys [32]) are another way to improve performance [34]; however, studies have shown that real-world use of these tools is often limited [43].

Accommodating a large number of commands within a selection technique is also an important issue, because typical ways of adding commands (e.g., with menus or ribbons) often add hierarchies which slow performance. Memory-based selection should allow a large command set while also maintaining fast access. A few examples exist for high-capacity techniques, such as Marking Menus [26] (64 items or more), ListMaps [20] (225 font items), CommandMaps [42] (210 items), or Kurtenbach et al.'s Hotbox, which supports large command sets by grouping the menu items into different Marking-Menu zones [27]. A problem for all large-capacity memory-based techniques, however, is that remembering command locations may become difficult as command set size increases.

Memory-based Interaction

Memory-based interfaces allow users to go directly to a command by recalling its location, rather than by visual-search-based navigation. Human memory is a well-studied topic, both in HCI and psychology (e.g., [15,17,38,46]). Numerous techniques such as gestures [28], hotkeys [32], spatial locations [11,21], or multi-touch chords [19]

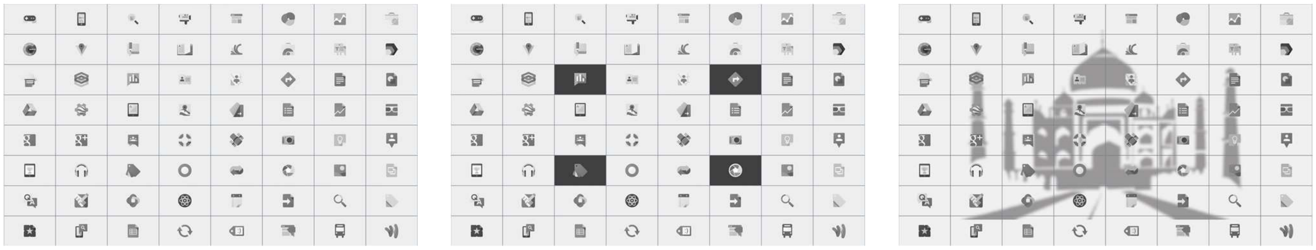

Figure 1. Interfaces with artificial landmarks (Study 1): (left) Standard menu with grid, (middle) Anchor menu with anchor points, and (right) Image menu with Taj Mahal's image as background.

demonstrated that people can build up extensive mappings between sets of items and command-invocation actions.

Gestures are a popular type of memory-based technique. For example, marking menus [26] and flower menus [6] provide a transition from navigation-based selection to memory-based gestures. In early use of these techniques, items can be found through visual search; but as users repeat selections for common items, they can begin to perform quicker selections by carrying out an accelerated and feedback-free version of the novice method. Other gesture techniques such as Octopocus [7] and Hotbox [27] try to aid the learning process by providing feedforward of possible gestures.

Past research with spatial memory in computer interfaces has shown that people can remember a large number of locations and can revisit them rapidly. For example, the Data Mountain [39] technique was significantly faster than ordinary bookmarking for retrieval of 100 web pages. Grid-based menus such as ListMap [20], FastTap [21], Square Menus [1], and CommandMaps [42] all showed performance advantages over either search-based or hierarchical organizations of data. However, there is still little understanding of the limits on spatial memory as a basis for user interfaces – and in particular, little understanding of how best to support location learning in these methods as command sets grow larger.

Use of Landmarks in Interfaces

In GUI-based systems, the landmarks that are already present in the environment (e.g., the corners of the screen) can provide a strong external reference frame that helps users build up spatial memory [47,48]. Several techniques have explicitly made use of the edges and corners of small devices (e.g., tablets or smartwatches) as landmarks to organize menus and toolbars [21,29,44]. However, these natural landmarks become less useful with larger screens, because many locations are not near a landmark. One technique for tabletops addressed this problem by using real-world objects (the user's own hands and fingers) to provide anchor points for faster location retrieval [49].

When natural landmarks are insufficient, artificially created visual elements can serve as landmarks [5,45]. Artificial landmarks, such as color [3] can help people to revisit an intended location quickly, and shape has also been used to give an object a memorable "visual ID" [30]. In addition, several techniques add marks to an interface to aid tasks such as understanding activity (e.g., Edit Wear and Read Wear [24]) and revisiting previously-seen items (e.g., Footprints Scrollbar [3] or Visual Popout UIs [18].

Several video summarization systems also use a type of landmark – e.g., creation of storyboards that indicate scene changes. For example, SceneSkim [35] provides browsing and skimming facilities using captions, scripts and plot summaries as reference points for different video locations, and Video Digests [36] represent sections/chapters with navigable markers. Other tools show visual highlights on timelines that represent personal [2] or crowd [25,52] navigation history, to support exploration and revisitation.

The design of landmarks has also been considered in 3D virtual environments to enhance first-person navigation [16,40], and guidelines exist for the design of landmarks for virtual worlds [50]. Less is known, however, about the design or value of landmarks in spatially-stable 2D interfaces. In our work, we use CommandMap's flat menu approach [42] to show a large number of commands in a grid-based overlay menu. We opted for this technique as it provides a basic representation on which people can develop spatial memory, and because it provides ample space for different types of artificial landmark.

LANDMARK TEST INTERFACES

To test the effect of different approaches to landmarking on spatial memory, we designed three similar interfaces (Figure 1) based on CommandMaps [42]. CommandMaps use all of the available display space to concurrently reveal all of the commands, each shown in a unique and stable spatial location. Normally, the CommandMap is not shown, allowing the full display space to be dedicated to the user's workspace (such as a document or spreadsheet). However, when the user wishes to access a command, they issue a control command (such as pressing a modifier key, mouse button, or gesture), which causes the CommandMap to be revealed. Command items are then selected by pointing and clicking on them. The CommandMap can be hidden either after each command selection, or on a subsequent control action (possibly allowing multiple commands to be invoked in a series, if the application requires it). As users learn the location of items in the CommandMap, they can anticipate the location at which they will be presented, facilitating rapid selection. Previous lab studies have demonstrated the efficiency of CommandMaps in comparison to menus and toolbar interfaces, both in abstract and realistic tasks [41].

As shown in Figure 1, our three landmark interfaces initially show a grid (*standard*, left), a grid plus a small set of dark gray grid anchors (*anchor,* middle), or a grid plus a transparent image of the Taj Mahal (*image*, right). When the user presses the Control key, a set of underlying icons are revealed in full screen setup. Selections are then made by clicking on the appropriate icon. Icons remain displayed until the Control key is released, or an icon item is clicked.

All three interfaces also support an *expert mode* of selection, in which icons can be selected prior to their display by pressing the Control key and immediately clicking in the location corresponding to the target item. To facilitate and encourage expert selections all interfaces implemented a timeout (200-400ms depending on the command set size) between pressing the Control key and displaying the icons.

We use the terms '*basic mode*' for selections that are completed with the aid of visual feedback after the short timeout, and '*expert mode*' for selections completed prior to the display of icons (Figure 2). As users' spatial memories of icon locations improve, they should complete more selections in *expert mode.* The landmarks used in the interfaces were always available in both *basic* and *expert mode* of selections (as shown in Figure 2).

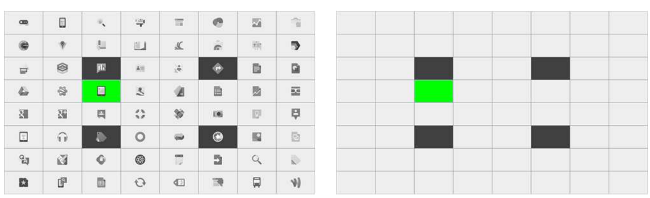

Figure 2. Selection modes: (left) basic - selection after icons are shown, (right) expert - selection without seeing the icons.

Standard Grid

The standard grid provides clear borders that demark item locations. In grids with relatively few items it is likely that the coarse placement resolution will allow the borders and corners to provide sufficient inherent landmark cues to assist spatial memorization (e.g., 'item by the top right corner').

Anchor

As the number of grid items increase, it is likely that the inherent landmark cues provided by borders and corners will become less effective in aiding spatial memorization and retrieval. Similarly, the grid lines provided by the *standard* condition are also likely to become less effective due to the frequency of their repetition.

The *anchor* condition therefore augments the spatial grid with dark grey grid cells that provide clear spatial reference points and are distinct from their surroundings [50]. Anchor cells are visible both before and after the icons are displayed. They therefore provide strong spatial anchors in *basic mode* selections, and particularly in *expert mode* where selections can be made without waiting for the icons to appear.

Image

The abstract landmarks provided in the *anchor* condition may provide inferior support for memory formation when compared to semantically meaningful objects. To test this possibility, the *image* condition overlays the grid with a simple greyscale image of a building (the Taj Mahal) with clear features (such as turrets, windows, doorways, paths, etc.). These features may assist memorization through the opportunities for association between objects (the icons) and the meaningful spatial location in the image (in a manner similar to the 'method of Loci' [37,51,54]). The inclusion of both an artificial landmark and image condition is partially motivated by prior findings on contextual cueing, which have indicated differences between item location learning in naturalistic scenes versus simple stimulus arrays [8,9]

STUDY 1: LANDMARKS IN A SMALL COMMAND SET

The following three studies examined user performance with our three landmark interfaces using progressively larger grids and with different lengths of training to enable memorization. Study 1 used an 8x8 grid of 64 items, and 11 blocks of trials. (Study 2 used a 8x12 grid (96 items) and 16 blocks of trials; Study 3 used a 10x16 grid and 18 blocks). All three studies were designed to answer the same main question: do the different approaches to landmarking result in different performance and styles of use?

Study 1 – Method

Tasks and stimulus. The study consisted of a series of trials, each involving the 'point and click' selection of a cued icon stimulus. Each trial began by displaying the stimulus icon on the left screen of a dual-monitor (21-inch) environment. The participant then used a mouse to select the target icon using one of the three landmarking interfaces (Figure 1) described above (*standard, anchor,* or *image*) which ran full-screen on the right screen.

The 64 icons (64px in size) presented in the interfaces were the same across all three interfaces (randomly relocated for each). They were extracted from the Android icon set [55] and converted to grayscale to reduce potential confounds from hue-induced popout effects. Twelve of the icons were quasi-randomly selected for use as stimuli with each interface condition. None of the icon stimuli were reused with a subsequent interface condition.

Procedure and study design. Participants were initially informed that the experiment concerned interfaces for rapid command selection. The *basic* and *expert* selection modes were described and demonstrated. They then completed 20 practice selections with each of the three interfaces using a different dataset to that used in the main experiment. Participants were instructed to complete trials as quickly and accurately as possible.

Participants completed 11 blocks of trials with each of the three interface conditions (order was counterbalanced). Having completed all blocks with one interface, participants completed a NASA-TLX [22] subjective workload questionnaire. They then progressed to the next interface.

Each of the first 10 blocks consisted of one trial for each of the 12 targets. Users were free to complete each trial using whichever selection modality they preferred (*basic* or *expert*). The 11th block involved a 'blind' trial for each target, with the *basic* selection mode disabled – participants clicked on the location they believed corresponded with the cued icon, without visual feedback.

Within each trial, a correct selection was confirmed by highlighting the selected grid location green for 400 ms; red for incorrect. Trials continued until correctly completed. Software recorded trial completion time, errors, expert selections, and data describing every selection. At the end of the study, participants chose their preferences for the three interfaces for various aspects of interaction.

Participants. Twelve participants (1 female), ages 18-34 (mean 25.5), were recruited from a local university. The study took ~60 minutes, and a $10 remuneration was paid to each participant.

Apparatus. The experiment was conducted on a desktop computer running Windows 7, with two 21-inch 1650x1050 resolution monitors placed alongside. Software was written in Java. Input was received through a standard keyboard and optical mouse. All study interfaces ran full-screen in the right of two 21-inch monitors.

Study 1 – Results

For all the three studies, we report the effect size for significant RM-ANOVA results as partial eta-squared: η^2 (considering .01 small, .06 medium, and >.14 large [7]). In all studies, where ANOVAs sphericity assumption is violated (Mauchley's test), Greenhouse-Geisser adjustments are performed (yielding floating point degrees of freedom).

Trial time

Mean trial completion times with the three interfaces are summarized across blocks in Figure 3.

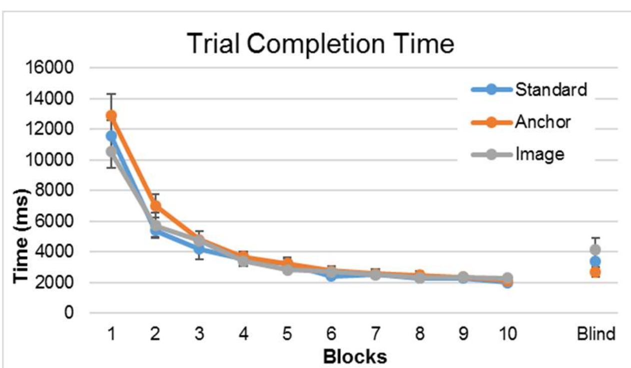

Figure 3. Mean trial time (±s.e.) by interface and block.

For the ten main blocks, RM-ANOVA showed no significant main effect of *interface* ($F_{2,22}$=2.0, p=.16), with means of 3940ms (s.d. 3146ms) with *standard*, 4348ms (s.d. 3727ms) with *anchor*, and 3931ms (s.d. 3003ms) with *image.*

Trial times decreased across *block* ($F_{1.84,20.2}$=78.84, p<.001, η^2=0.88), and as anticipated, the skill development follows a power law function [33]. There was no significant *interface* × *block* interaction ($F_{18,198}$=1.04, p=.42).

Analysis of mean trial time in the final blind block also showed no significant difference between the three interfaces ($F_{2,22}$=2.1, p=.15).

Error rates and expert selections

Analysis of the number of errors per trial showed similar results to the trial time analysis. There was no main effect of *interface* ($F_{2,22}$ < 1) and no *interface* × *block* interaction ($F_{18,198}$ < 1). Errors increased with *block* ($F_{9,99}$=2.94, p=.004, η^2=0.21), which can be attributed to two potential causes – users becoming faster and less precise, and users increasingly attempting to rely on incompletely-formed spatial memories with the expert mode.

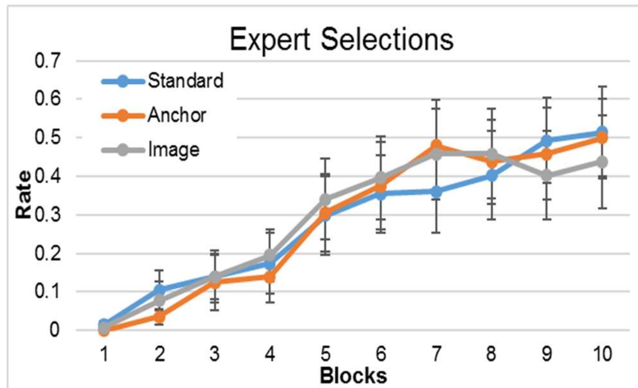

Figure 4. Expert selection rates by interface and block.

Use of the expert selection mode followed a similar pattern to errors, with only *block* ($F_{1.63,17.92}$=14.96, p<.001, η^2=0.58) showing a significant effect (Figure 4). Participants made the same proportion of expert selections with all three interfaces (overall, 0.29 selections/trial).

Subjective responses

Participant responses on the NASA-TLX worksheets were also similar for the three interfaces (see Table 1), with no significant effects except for reported Frustration: *image* (mean 3.25, s.d. 2.3) induced higher frustration than *anchor* (2.8, s.d. 2.7) and *standard* (2.3, s.d. 1.2).

	Standard	Anchor	Image	χ_r^2	p
Mental	7.17(1.67)	6.58(1.8)	6.42(1.38)	1.79	.41
Physical	2.67(2.49)	2.67(2.32)	3.42(2.6)	1.54	.46
Temporal	5.25(2.05)	5.75(1.69)	5.5(2.06)	0.79	.67
Performance	5.00(2.74)	5.33(2.95)	4.25(2.45)	3.16	.21
Effort	6.00(2.31)	6.00(1.91)	5.92(1.61)	0.13	.94
Frustration	2.25(1.23)	2.75(2.68)	3.25(2.28)	3.12	**.01**

Table 1. Mean (s.d.) effort scores (0-10 scale, low to high).

	Standard	Anchor	Image
Speed	1	8	3
Accuracy	1	8	3
Memorization	0	8	4
Expertise	3	7	2
Comfort	1	9	2
Overall	1	8	3

Table 2. Count of participant preferences.

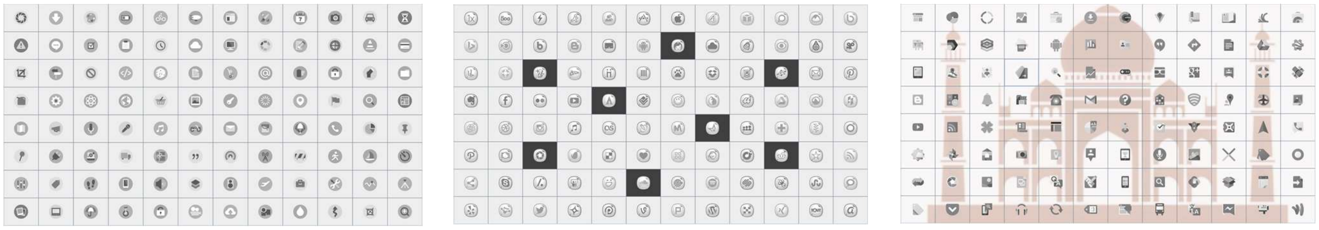

Figure 5. Study interfaces for Study 2: (left) Standard menu with grid, (middle) Anchor menu with anchor points, and (right) Image menu with Taj Mahal's image as background.

Despite the lack of objective and subjective workload findings favoring any interface, the participants' preferences were strongly in favor of the *anchor* interface. Eight or more of the twelve participants selected it as the preferred interface for Speed, Accuracy, Memorization, Comfort and Overall (see Table 2).

STUDY 2: LANDMARKS IN A MEDIUM COMMAND SET

The second study used a similar method to Study 1, but with the following alterations. First, the grid used in all three interfaces was larger, with 96 items (64px) in a 8x12 arrangement; the item set was unique in each condition to avoid learning effects. Second, there were 15 blocks of trials instead of 10; and a final 16th 'blind' block was used in the same manner as Study 1. Third, nine targets were used in each block rather than 12; this adjustment, together with the increase in the number of blocks allows greater opportunity for spatial learning with all interfaces. Fourth, in the *anchor* condition there were 8 dark gray items rather than 4, positioned as shown in Figure 5. Last, we increased the timeout delay to 350ms from 200ms. These adjustments were made to accommodate the higher number of items.

In other aspects, the method, procedure, and apparatus was identical to Study 1.

Participants. Twelve participants (5 females), aged 19-37 (mean 27) were recruited at a local university. None had previously participated in Study 1. The study lasted for approximately 60 minutes, and a $10 honorarium was paid to each participant.

Study 2 – Results

Trial time

Mean trial times with the three interfaces across blocks are shown in Figure 6. RM-ANOVA showed no significant main effect of *interface* ($F_{2,22}$=3.38, p=.05). Mean trial times were lowest with *anchor* (3990ms, s.d. 4643ms), followed by *standard* (5057ms, s.d. 4464ms) and *image* (7463ms, s.d. 4204ms).

As in Study 1, there was a significant effect of *block* ($F_{2.35,25.87}$=64.1, p<.001, η^2=0.85), but no *interface* × *block* interaction ($F_{28,308}$=1.1, p=.34).

Unlike Study 1, analysis of mean trial time in the final *blind* block showed a significant difference between the three *interfaces* ($F_{1.13,12.45}$=3.75, p=.07, η^2=0.25), with *anchor* being the fastest (mean 1717ms, s.d. 309ms).

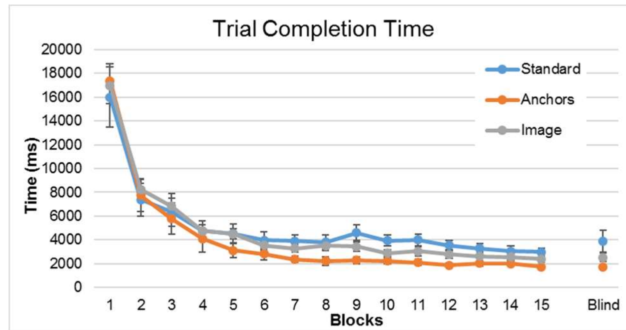

Figure 6. Mean trial completion time by interface and block.

Error rates

There was a significant main effect of *interface* on errors ($F_{2,22}$=9.2, p=.001, η^2=0.46). *Anchor* had the lowest error rate at 0.07 errors/trial (s.d. 0.13), compared to much higher error rates of 0.23 (0.25) and 0.27 (0.3) errors/trial with *standard* and *image* respectively. Bonferroni-corrected follow-up t-tests showed that *anchor* was more accurate than both of the other interfaces (all p<.001).

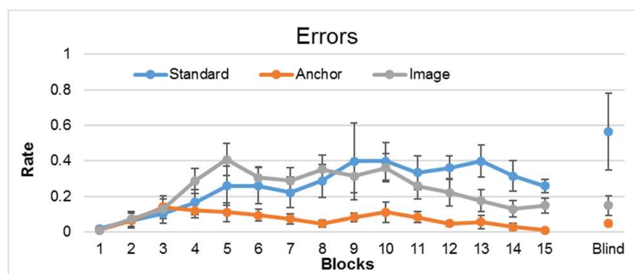

Figure 7. Error rate across blocks with the three interfaces.

As shown in Figure 7, with all of the interfaces, errors increased across the first blocks, then roughly stabilized, before decreasing in the final blocks, leading to a significant effect of *block* ($F_{4.09,44.97}$=5.99, p<.001, η^2=0.35). There was a significant *interface* × *block* interaction ($F_{28,308}$=2.17, p=.001, η^2=0.16), attributable to *Anchor* having comparatively stable and low errors across blocks.

We also analsyed error rates based on the location of target items in the grid (along the side, at the corner, or in the middle), and we categorized errors based on the distance from the intended target ('off by 1' and 'off by more than 1'). Table 3 summarizes the findings, showing the proportion of trials at each location that contained an error at each distance. The table reveals some interesting additional characteristics of errors. With *standard* and *image* the highest error rates occurred in the middle of the grid. With *standard*, the total error rate for middle targets was 34.3% (23.7+10.6), and with *image* it was 31.5%. These high rates contrast with the relatively low value of 8.7% for middle targets with *anchors.* Corner error rates were much lower with all three interfaces (*standard* 11.4, *anchors* 4.7, and *image* 6.7%), which can be explained by the unambiguous spatial demarcation provided by the corner. Interestingly, side errors were also much lower with *anchors* (5.8%) than *standard* (18.0%) and *image* (18.6%), which may be due to the *anchors* providing clear reference points along both independent dimensions (e.g., "on the right edge and aligned with the top gray block")

	Standard		Anchor		Image	
Group	Off by 1	Off by 1+	Off by 1	Off by 1+	Off by 1	Off by 1+
Side	14.4	3.6	**3.9**	1.9	12.8	5.8
Corner	7.5	3.9	**2.8**	1.9	5.0	1.7
Middle	23.7	10.6	**2.1**	6.6	21.7	9.8

Table 3. Percentage of errors at each target location group (side, corner or middle), broken down by error distance.

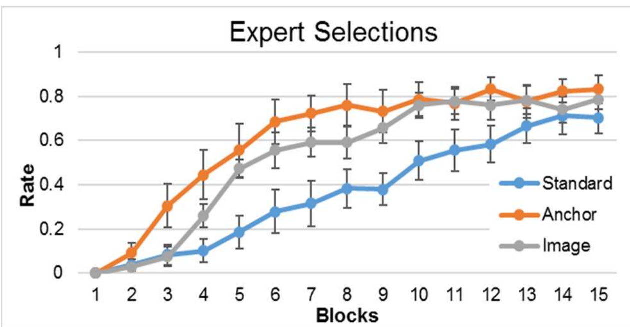

Figure 8. Expert selection rates by interface and block.

Use of expert selections

Figure 8 shows the rate of expert mode selections with the three interfaces across block. It shows that the *standard* interface had much lower use of the expert mode than the other interfaces in nearly all blocks (mean 0.37 sel./trial, s.d. 0.35), and that the *anchor* interface had the highest rate (mean 0.61 selections/trial, s.d. 0.38); expert mode selections with *image* were slightly lower (0.52 sel./trial, s.d. 0.34): $F_{2,22}$=6.74, p=.005, η^2=0.38). There was a significant *interface* × *block* interaction ($F_{28,308}$=2.27, p<.001, η^2=0.17), as indicated in Figure 8. Post-hoc t-tests (Bonferroni corrected) show that *anchor* performed significantly better than the other two interfaces, and that *image* was also faster than *standard* (all p<.001).

Subjective responses

Table 4 summarizes mean response to the NASA-TLX worksheets. Friedman tests showed significant effects for Mental workload (with *image* lowest, followed by *anchor* and *standard*), Temporal workload (*anchor* lowest, then *image* and *standard* highest) and Performance (*anchor* highest*, standard* lowest).

	Standard	Anchor	Image	χ_r^2	p
Mental	7.17(2.43)	5.42(2.33)	4.33(2.17)	9.04	**.01**
Physical	3.17(3.21)	2.92(2.93)	2.75(2.77)	0.29	.86
Temporal	5.42(2.75)	4.00(2.86)	5.08(2.72)	6.5	**.04**
Performance	5.75(1.69)	7.25(1.3)	6.00(1.96)	10.79	**.01**
Effort	6.83(1.91)	5.33(2.87)	6.42(2.14)	5.79	.06
Frustration	4.00(2.58)	2.75(2.42)	3.33(3.06)	3.79	.15

Table 4. Mean (s.d.) effort scores (0-10 scale, low to high).

Counts of the preferred interface strongly favored the *anchor* interface, with 83-92% of participants selecting it as preferred across six dimensions (see Table 5).

	Standard	Anchor	Image	None
Speed	0	10	2	0
Accuracy	0	10	2	0
Memorization	0	11	1	0
Expertise	0	10	1	1
Comfort	0	11	1	0
Overall	0	11	1	0

Table 5. Count of participant preferences.

STUDY 3: LANDMARKS IN A LARGE COMMAND SET

The third study used the method of Study 2, but with a grid of 160 items (48px) in a 10x16 grid arrangement. Additionally, to provide stronger insights into the formation of spatial memory, we used a blind block of trials after every 5th block. There were therefore 18 blocks in total, consisting of three repetitions of 5 regular blocks followed by one blind block. Ten targets were used instead of the nine in Study 2. Also, we increased the number of gray blocks in the *anchor* condition from 8 to 14 (positions are shown in Figure 9). Finally, we set the timeout delay for expert mode selection to 400ms. The method, procedure and apparatus were identical to Study 2 in other aspects.

Participants. We recruited 16 new participants (9 females), ages 19-41 (mean 28.4), from a local university. The study lasted ~60 minutes. Participants received a $10 payment.

Study 3 – Results

Trial time

Mean trial times for the 15 *basic* blocks and 3 'blind' blocks (analyzed separately) are shown in Figure 10. RM-ANOVA showed a significant main effect of *interface* on completion

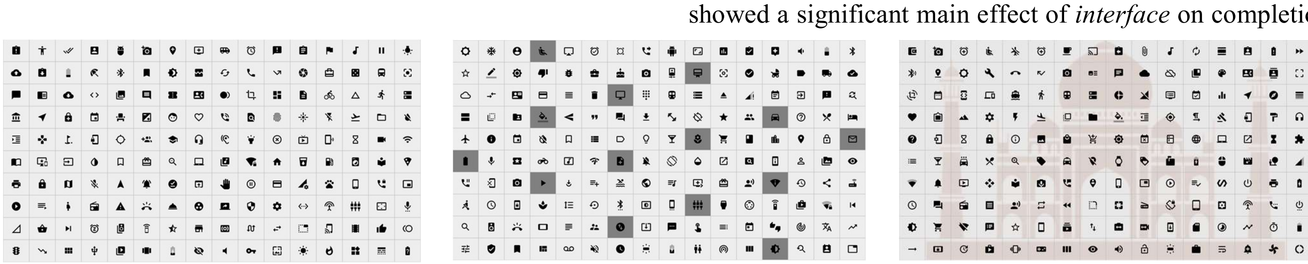

Figure 9. Study interfaces for Study 3: (left) Standard menu with grid, (middle) Anchor menu with anchor points, and (right) Image menu with Taj Mahal's image as background.

time ($F_{2,30}$=4.4, p=.02, η^2=0.23). In the *basic* blocks, *anchor* was the fastest with mean 5108ms (s.d. 5620ms) compared to *image* (5215ms, s.d. 4515ms) and *standard* (6416ms, s.d. 6196ms). Bonferroni-corrected t-tests showed that both *anchor* and *image* were significantly faster than *standard* (p<.001), but there was no difference between *anchor* and *image* (p=.68).

Trial time significantly decreased across *block* ($F_{2.1,31.5}$=66.8, p<.001, η^2=0.82), and there was also a significant *interface* × *block* interaction ($F_{28,420}$=2.5, p<.001, η^2=0.14). For the three *blind* blocks (see Figure 10) there was no main effect of *interface*, ($F_{2,30}$=2.5, p=.1), but there was a significant effect of *block* ($F_{1.19,17.91}$=9.98, p=.004, η^2=0.4), with mean times becoming slightly faster in later *blind* blocks. There was no *block* × *interface* interaction.

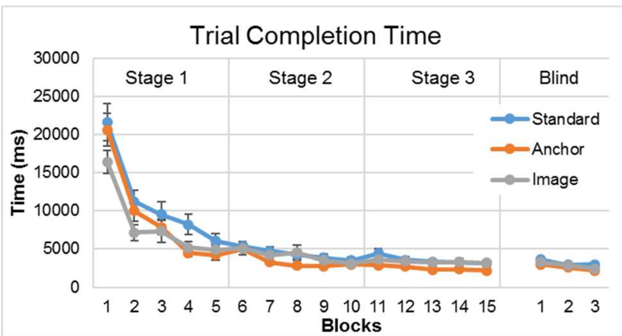

Figure 10. Mean trial completion time by interface and block.

Error rates

For the 15 *basic* blocks, there was no significant main effect of *interface* (see Figure 11): $F_{1.37,20.52}$=1.52, p=.24. The lack of a significant effect of *interface* differs from Study 2, and is probably due to the additional difficulty of learning 10 items among 160 (rather than 9 among 96) causing more random variation in Study 3. There was a significant effect of *block* ($F_{4.56,68.46}$=3.53, p=.008, η^2=0.19,), but no *interface* × *block* interaction ($F_{28,420}$=0.84, p=.71).

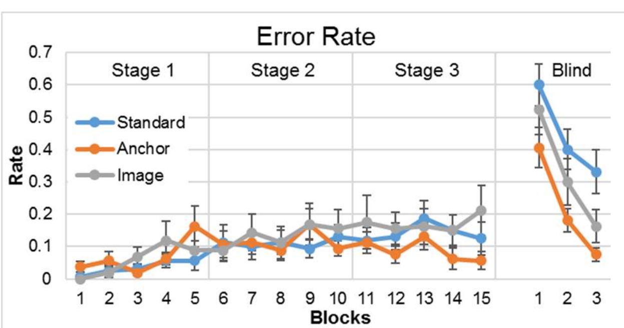

Figure 11. Mean error rates in Study 3.

For the *blind* blocks, *anchor* had the lowest error rate (0.22 errors/trial, s.d. 0.22), followed by *image* (0.33, s.d. 0.31) and standard (0.44, s.d. 0.28), giving a significant effect of *interface* ($F_{2,30}$=9.88, p=.001, η^2=0.4). Errors significantly decreased across the three blocks ($F_{2.1,31.5}$=36.18, p<.001, η^2=0.71), with participants making far fewer errors as they learned spatial locations. There was no *interface* × *block* interaction.

Analysis of the influence of target locations (side, corner and middle) on the proportion and distance of errors revealed similar observations to those reported in Study 2 (see Table 6). For all interfaces, the highest proportion of errors occurred with *middle* targets. With the *anchor* interface, *middle* errors were lower (14.6%) than *standard* (19.4%) and *image* (20.2%) interfaces. *Corner* errors were similar across interfaces (4.3, 6.7, and 6.4 with *standard, anchor*, and *image*). *Side* errors were also lower with *anchors* (5.2%) than with *standard* (14.4%) and *image* (10.9%).

	Standard		Anchor		Image	
Group	**Off by 1**	**Off by 1+**	**Off by 1**	**Off by 1+**	**Off by 1**	**Off by 1+**
Side	9.9	4.5	**3.5**	1.7	6.6	4.3
Corner	4.0	0.3	3.1	3.6	5.4	1.0
Middle	13.1	6.3	**3.1**	11.5	10.5	9.7

Table 6. Percentage of errors at each target location group (side, corner or middle), broken down by error distance.

Use of expert selections

As shown in Figure 12, during *basic* blocks, *anchor* had the highest level of expert selections (mean 0.44 sel./trial, s.d. 0.39) compared to *image* (0.39, s.d. 0.4), and *standard* (0.26, s.d. 0.34): $F_{2,30}$=7.39, p=.002, η^2=0.33. In the final blocks of Stage 3 (when users had had the greatest opportunity to learn locations), approximately 80% of *anchor* selections were made using the expert modality, compared to 65% with *image* and 50% with *standard*. Post-hoc t-tests (Bonferroni corrected) showed differences in expert selection rate for all *interface* pairs (all p<.001).

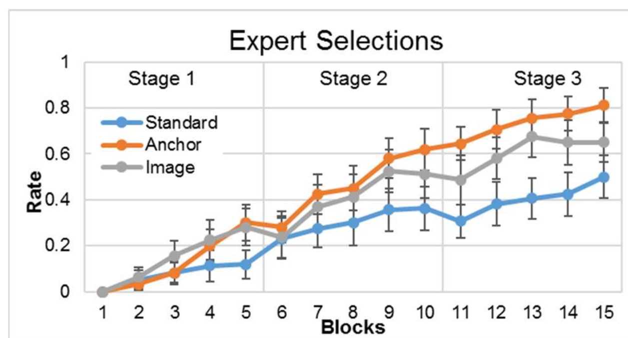

Figure 12. Expert selection rates by interface and block.

Subjective responses

NASA-TLX responses were also similar to Study 2 (see Table 7). *Anchor* and *image* interfaces received lower workload scores and higher performance scores. Preference counts also favored the *anchor* interface, with no participants selecting *standard* as preferred (see Table 8).

	Standard	Anchor	Image	χ_r^2	p
Mental	7.38(1.73)	6.38(2.06)	6.13(2.18)	7.63	**.02**
Physical	2.69(2.49)	2.5(2.06)	2.31(2.02)	0.38	.83
Temporal	5.75(2.68)	4.94(2.49)	4.56(2.6)	6.84	**.03**
Performance	4.44(1.73)	7.19(2.01)	6.75(1.92)	14.28	**<.01**
Effort	7.00(2.18)	6.00(1.9)	6.06(1.52)	5.34	.07
Frustration	5.94(2.9)	4.81(2.74)	4.94(2.84)	4.72	.09

Table 7. Mean (s.d.) effort scores (0-10 scale, low to high).

	Standard	Anchor	Image	None
Speed	0	8	7	1
Accuracy	0	8	6	2
Memorization	0	8	7	1
Expertise	0	9	4	3
Comfort	0	8	6	2
Overall	0	8	6	2

Table 8. Count of participant preferences.

Participant Comments

Participant comments for the three studies mirrored and emphasized the objective findings. Participants made several comments on how the artificially planted landmarks (especially *anchors*) helped them to develop spatial memory of the commands and improve performance: one participant mentioned *"[Target] locations were easy to predict because of the anchors [in the grid]."* Another said *"It was easy to remember 5-7 anchors and it helped me to find other [target] items."* One person, however, remarked on the difficulty of remembering commands in the *standard* condition: *"Too many items [in grid with no landmark] made it hard to remember the position, although I tried to use grid number to remember [targets]."*

Other comments suggested that the *image* also helped participants to learn command locations: one said *"It was easier to remember things when there was meaningful content [in image] to connect it to (e.g. reading a book on the second floor)."* Another observed that *"I associated the objects with various parts of the monument [image], as if the library was located in one of the minarets of the Taj Mahal."*

However, landmarks can cause distraction and often require extra memory to process information. One person stated *"[Often] when I started looking for any icon, [the] background image caught my attention."*

DISCUSSION

The main findings of the three studies are as follows:

- *Anchor* landmarks improved users' ability to memorize grid item locations, compared to an unadorned grid. By providing abstract spatial cues, users were able to more quickly and accurately select target items.
- In a relatively small grid of 64 items, there were no significant differences between performance with the *standard, anchor*, and *image* interfaces. However, preferences strongly favored the *anchor* interface. It appears that inherent spatial cues provided by the *standard* interface (including the grid, display edge and corners) were sufficient to support effective memorization.
- However, in larger grids, the additional landmarking features provided by the *anchor* and *image* interfaces enabled better memorization and retrieval. Analysis of the proportion of errors by target location (middle, edge, or corner of the display) indicated that errors with the *standard* interface were most prevalent when spatial landmarks were less clear – errors were lowest at the spatially unambiguous display corners, higher along the edges, and highest in the middle of the display (away from corners and edges).
- Mean selection times and error rates were lower with the abstract spatial landmarks provided by *anchor* than they were with *image.* This is interesting because we initially suspected that the *image* interface might present more opportunities for semantic association (e.g., 'the star is by the left minaret), but this was not the case. In addition, *anchor* was strongly preferred over *image*.

All interfaces supported spatial learning

The reduction in mean trial completion time across blocks conforms to the expected power-law of learning curves [33] in all three studies. This can be attributed to users transitioning from selections that are dominated by slow visual search (followed by rapid pointing) to selections that are characterized by rapid spatial recollection and pointing. This transition occurs with all three interfaces, indicating that users developed spatial memories, regardless of interface. However, Figures 8 and 12 from studies 2 and 3 show that participants were able to form and exploit these memories more rapidly when additional landmarking features were available in *anchor* and *image* interfaces.

Why did *anchor* outperform *image?*

Trial time data, error rates, workload measures and preferences all favored the *anchor* interface over *image.* Yet centuries of evidence from the 'method of loci' [37,54] suggest that concrete spatial representations (such as buildings) can assist memorization of abstract concepts by associating those concepts with the spatial representation.

We see three main reasons for the comparatively strong performance of the *anchor* interface compared to *image*. First, like many user interfaces, the items used in our study were presented aligned to a clear two dimensional grid. The gray blocks of the *anchor* condition therefore provided strong alignment cues on each dimension – even when a target was distant from one of the blocks, it could serve as a spatial cue ("same row as that gray block over there"). In contrast, the Taj Mahal image less clearly afforded this form of dimension-based alignment – participants may have been less likely to exploit independent row and column alignment concepts when targets were distant from image features.

Second, the gray blocks of the *anchor* interface had high visual salience – users were almost certain to incorporate these landmarks into their conception of the tasks (e.g., "the star is the black top-left block", or "the star is the item above the black top-left block"). In contrast, the Taj Mahal image was more subtly presented, and it might have been ignored by participants during their tasks. Indeed, one participant commented *"I did not notice the background image in grid."* In addition, the finer-grained details on the image may have been more difficult for users to incorporate in their memory, since there was more to remember about the landmark (e.g., "near to this particular filigree on the roofline of the building").

The *anchor* and *image* conditions can be viewed as representing different points on two continuums between sparse and dense image features, and between abstract and concrete representations. Further study is needed to examine the role of these visual characteristics on spatial memorization in user interfaces – but our study suggests that there could be value in simpler representations when providing landmarks that are adjunct to the primary task.

Third, although the Taj Mahal is a well-known building, none of our participants had personal experience of it. The method of loci, however, is based on the intentional placement of memorization objects into a highly personal environment (such as a favorite walk or the rooms of one's home). Results might differ with a more personally-familiar image.

Implications for design

The findings suggest that users' ability to form and draw on spatial memories for rapid interaction will be assisted by the presentation of landmarks. Many interface components that rely on spatial interaction are often featureless, blank spaces, creating problems for users in exploiting their spatial memories. One example is the blank trough of a scrollbar, which can give rise to inefficient interactions. Previous researchers have proposed augmenting the scroll-trough with transient markers to indicate recently or frequently visited regions, which assist users returning to their place in a document after cross-referral to another area [3,24]. Our results suggest that static embellishments in the scroll-trough could provide useful landmarks for associating spatial locations.

Of course, there are challenges to the use of landmarks in user interfaces. The first of these involves the potential for interference from the artificial landmarks (which are always present on the screen) and the user's document content. If the user is working in a graph editor that shows rectangular blocks on the screen, for example, the anchor visualizations could conflict with task objects. This could result in interference in both directions: the artificial landmarks could hinder the interpretation of the document contents, and document graphics could potentially interfere with the value of the anchor landmarks. We believe that this issue can be readily addressed through careful design of the artificial landmarks. For example, the landmarks can be extremely faint (e.g., using a very high level of transparency) and still be useful – once the user is familiar with them, they will need only a minimal representation to guide their spatial memory. In addition, the visual representation of anchor landmarks can potentially be changed without affecting the user's memory – for example, specific colors or textures could be chosen so that they do not interfere with the visual features of the document content.

A related challenge in the use of artificial landmarks involves aesthetics – that is, there is a potential impairment of design aesthetics due to the presentation of otherwise superfluous visual features, and the risk that users will interpret the landmarks as being part of the design of the application rather than as spatial reference points. Further work, particularly in collaboration with graphic designers, is needed to address these concerns.

Although *image* provided richer landmark features than *anchor*, surprisingly it did not perform as we expected. The Taj Mahal images we used in our studies were carefully converted into grayscale and later faded out to avoid color overlap with any task icons. It is possible, however, that the background image may have still interfered with a few of the icons – we will consider this issue in further item-by-item analyses. However, *image* provided an overall performance advantage beyond that attained with the *standard* grids. It will be an interesting to investigate the effect of different feature rich images as artificial landmarks.

In future work we plan to carry out further studies of the value of artificial landmarks, in practical settings with realistic document content, and in different types of applications. For example, the ideas presented here can be tested in one-dimensional representations such as scroll bars and video timelines, as well as in two-dimensional settings such as desktop/homescreen wallpaper decorations. We will also carry out new studies to explore potential interference between artificial landmarks and other objects on the screen, and we will test new designs to determine whether subtle landmark representations that are visually unobtrusive can still provide effective anchors for spatial memory.

CONCLUSION

Desktop and mobile user interfaces make heavy use of spatial organization to facilitate rapid access to interface items. We examined the role that landmarks play in assisting spatial memorization and retrieval of items in a grid of interface components. The landmarks were static and passive visual embellishments designed to help users orient themselves in the graphical layout. Three forms of landmarking assistance were empirically compared – basic gridlines, the additional use of grey fill for some grid cells to provide clear visual anchors, and the grid overlaid with a meaningful background image. Item retrieval times and error rates were best when using the simple visual anchors. Reasons for the findings, implications for design, and directions for further work were presented.

ACKNOWLEDGMENTS

This work was supported by Natural Sciences and Engineering Research Council of Canada.

REFERENCES

1. David Ahlström, Andy Cockburn, Carl Gutwin, and Pourang Irani. 2010. Why it's quick to be square: modelling new and existing hierarchical menu designs. *Proceedings of the 28th international conference on Human factors in computing systems - CHI '10*, ACM Press, 1371–1380. http://doi.org/10.1145/1753326.1753534
2. Abir Al-Hajri, Gregor Miller, Matthew Fong, and Sidney S. Fels. 2014. Visualization of personal history for video navigation. *Proceedings of the 32nd annual*

ACM conference on Human factors in computing systems - CHI '14, ACM Press, 1187–1196. http://doi.org/10.1145/2556288.2557106

3. Jason Alexander, Andy Cockburn, Stephen Fitchett, Carl Gutwin, and Saul Greenberg. 2009. Revisiting read wear: analysis, design, and evaluation of a footprints scrollbar. *Proceedings of the 27th international conference on Human factors in computing systems - CHI '09*, ACM Press, 1665–1674. http://doi.org/10.1145/1518701.1518957
4. Fraser Anderson and Walter F. Bischof. 2013. Learning and performance with gesture guides. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, ACM Press, 1109–1118. http://doi.org/10.1145/2470654.2466143
5. Wendy Ark, D Christopher Dryer, Ted Selker, and Shumin Zhai. 1998. Landmarks to aid navigation in a graphical user interface. *Workshop on Personalised and Social Navigation in Information Space*. Retrieved September 15, 2016 from http://www.shuminzhai.com/papers/sweden/final.html
6. Gilles Bailly, Eric Lecolinet, and Laurence Nigay. 2008. Flower menus: a new type of marking menu with large menu breadth, within groups and efficient expert mode memorization. *Proceedings of the working conference on Advanced visual interfaces - AVI '08*, ACM Press, 15–22. http://doi.org/10.1145/1385569.1385575
7. Olivier Bau and Wendy E. Mackay. 2008. OctoPocus: a dynamic guide for learning gesture-based command sets. *Proceedings of the 21st annual ACM symposium on User interface software and technology - UIST '08*, ACM Press, 37–46. http://doi.org/10.1145/1449715.1449724
8. James R. Brockmole, Monica S. Castelhano, and John M. Henderson. 2006. Contextual cueing in naturalistic scenes: Global and local contexts. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 32, 4: 699–706. http://doi.org/10.1037/0278-7393.32.4.699
9. James R. Brockmole and John M. Henderson. 2006. Using real-world scenes as contextual cues for search. *Visual Cognition* 13, 1: 99–108. http://doi.org/10.1080/13506280500165188
10. John R. Callahan, Don Hopkins, Mark D. Weiser, and Ben Shneiderman. 1988. An empirical comparison of pie vs. linear menus. *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '88*, ACM Press, 95–100. http://doi.org/10.1145/57167.57182
11. Andy Cockburn, Carl Gutwin, and Jason Alexander. 2006. Faster document navigation with space-filling thumbnails. *Proceedings of the SIGCHI conference on Human Factors in computing systems - CHI '06*, ACM Press, 1–10. http://doi.org/10.1145/1124772.1124774
12. Andy Cockburn, Carl Gutwin, and Saul Greenberg. 2007. A predictive model of menu performance. *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07*, ACM Press, 627–636. http://doi.org/10.1145/1240624.1240723
13. Andy Cockburn, Per Ola Kristensson, Jason Alexander, and Shumin Zhai. 2007. Hard lessons: effort-inducing interfaces benefit spatial learning. *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '07*, ACM Press, 1571–1580. http://doi.org/10.1145/1240624.1240863
14. Fergus I.M. Craik and Robert S. Lockhart. 1972. Levels of processing: A framework for memory research. *Journal of Verbal Learning and Verbal Behavior* 11, 6: 671–684. http://doi.org/10.1016/S0022-5371(72)80001-X
15. Mary Czerwinski and George Robertson. 1999. The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D. *Proceedings of Seventh International Conference on Human-Computer Interaction - INTERACT '99*, IOS Press, 163–170.
16. Rudolph P. Darken and John L. Sibert. 1996. Wayfinding strategies and behaviors in large virtual worlds. *Proceedings of the SIGCHI conference on Human factors in computing systems common ground - CHI '96*, ACM Press, 142–149. http://doi.org/10.1145/238386.238459
17. Brian D. Ehret. 2002. Learning where to look: location learning in graphical user interfaces. *Proceedings of the SIGCHI conference on Human factors in computing systems Changing our world, changing ourselves - CHI '02*, ACM Press, 211–218. http://doi.org/10.1145/503376.503414
18. Krzysztof Z. Gajos, Mary Czerwinski, Desney S. Tan, and Daniel S. Weld. 2006. Exploring the design space for adaptive graphical user interfaces. *Proceedings of the working conference on Advanced visual interfaces - AVI '06*, ACM Press, 201–208. http://doi.org/10.1145/1133265.1133306
19. Emilien Ghomi, Stéphane Huot, Olivier Bau, Michel Beaudouin-Lafon, and Wendy E. Mackay. 2013. Arpège: learning multitouch chord gestures vocabularies. *Proceedings of the 2013 ACM international conference on Interactive tabletops and surfaces - ITS '13*, ACM Press, 209–218. http://doi.org/10.1145/2512349.2512795
20. Carl Gutwin and Andy Cockburn. 2006. Improving list revisitation with ListMaps. *Proceedings of the working conference on Advanced visual interfaces - AVI '06*, ACM Press, 396–403. http://doi.org/10.1145/1133265.1133347

21. Carl Gutwin, Andy Cockburn, Joey Scarr, Sylvain Malacria, and Scott C. Olson. 2014. Faster command selection on tablets with FastTap. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, ACM Press, 2617–2626. http://doi.org/10.1145/2556288.2557136
22. Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Advances in Psychology* 52: 139–183. http://doi.org/10.1016/S0166-4115(08)62386-9
23. William E. Hick. 1952. On the rate of gain of information. *Quarterly Journal of Experimental Psychology* 4, 1: 11–26. http://doi.org/10.1080/17470215208416600
24. William C. Hill, James D. Hollan, Dave Wroblewski, and Tim McCandless. 1992. Edit wear and read wear. *Proceedings of the SIGCHI conference on Human factors in computing systems - CHI '92*, ACM Press, 3–9. http://doi.org/10.1145/142750.142751
25. Juho Kim, Philip J. Guo, Carrie J. Cai, Shang-Wen (Daniel) Li, Krzysztof Z. Gajos, and Robert C. Miller. 2014. Data-driven interaction techniques for improving navigation of educational videos. *Proceedings of the 27th annual ACM symposium on User interface software and technology - UIST '14*, ACM Press, 563–572. http://doi.org/10.1145/2642918.2647389
26. Gordon Kurtenbach and William Buxton. 1994. User learning and performance with marking menus. *Proceedings of the SIGCHI conference on Human factors in computing systems celebrating interdependence - CHI '94*, ACM Press, 258–264. http://doi.org/10.1145/191666.191759
27. Gordon Kurtenbach, George W. Fitzmaurice, Russell N. Owen, and Thomas Baudel. 1999. The Hotbox: efficient access to a large number of menu-items. *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*, ACM Press, 231–237. http://doi.org/10.1145/302979.303047
28. Gordon Kurtenbach, Abigail Sellen, and William Buxton. 1993. An Empirical Evaluation of Some Articulatory and Cognitive Aspects of Marking Menus. *Human-Computer Interaction* 8, 1: 1–23. http://doi.org/10.1207/s15327051hci0801_1
29. Benjamin Lafreniere, Carl Gutwin, Andy Cockburn, and Tovi Grossman. 2016. Faster Command Selection on Touchscreen Watches. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, ACM Press, 4663–4674. http://doi.org/10.1145/2858036.2858166
30. John P. Lewis, Ruth Rosenholtz, Nickson Fong, and Ulrich Neumann. 2004. VisualIDs: automatic distinctive icons for desktop interfaces. *ACM Transactions on Graphics* 23, 3: 416–423. http://doi.org/10.1145/1015706.1015739
31. Kevin Lynch. 1960. *The image of the city*. MIT Press.
32. Sylvain Malacria, Gilles Bailly, Joel Harrison, Andy Cockburn, and Carl Gutwin. 2013. Promoting Hotkey use through rehearsal with ExposeHK. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems - CHI '13*, ACM Press, 573–582. http://doi.org/10.1145/2470654.2470735
33. Allen Newell and Paul S. Rosenbloom. 1981. Mechanisms of skill acquisition and the law of practice. *Cognitive skills and their acquisition* 1: 1–55.
34. Daniel L. Odell, Richard C. Davis, Andrew Smith, and Paul K. Wright. 2004. Toolglasses, marking menus, and hotkeys: a comparison of one and two-handed command selection techniques. *Proceedings of Graphics Interface 2004 - GI '04*, Canadian Human-Computer Communications Society, School of Computer Science, Univ. of Waterloo, 17–24. Retrieved September 8, 2016 from http://dl.acm.org/citation.cfm?id=1006058.1006061
35. Amy Pavel, Dan B. Goldman, Björn Hartmann, and Maneesh Agrawala. 2015. SceneSkim: Searching and Browsing Movies Using Synchronized Captions, Scripts and Plot Summaries. *Proceedings of the 28th Annual ACM Symposium on User Interface Software & Technology - UIST '15*, ACM Press, 181–190. http://doi.org/10.1145/2807442.2807502
36. Amy Pavel, Colorado Reed, Björn Hartmann, and Maneesh Agrawala. 2014. Video Digests: a browsable, skimmable format for informational lecture videos. *Proceedings of the 27th annual ACM symposium on User interface software and technology - UIST '14*, ACM Press, 573–582. http://doi.org/10.1145/2642918.2647400
37. Simon T. Perrault, Eric Lecolinet, Yoann Pascal Bourse, Shengdong Zhao, and Yves Guiard. 2015. Physical Loci: leveraging spatial, object and semantic memory for command selection. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems - CHI '15*, ACM Press, 299–308. http://doi.org/10.1145/2702123.2702126
38. Albert Postma and Edward H.F. De Haan. 1996. What was where? Memory for object locations. *The Quarterly journal of experimental psychology. A, Human experimental psychology* 49, 1: 178–99. http://doi.org/10.1080/713755605
39. George Robertson, Mary Czerwinski, Kevin Larson, Daniel C. Robbins, David Thiel, and Maarten van Dantzich. 1998. Data mountain: using spatial memory for document management. *Proceedings of the 11th annual ACM symposium on User interface software and technology - UIST '98*, ACM Press, 153–162. http://doi.org/10.1145/288392.288596

40. Roy A. Ruddle, Ekaterina Volkova, and Heinrich H. Bülthoff. 2013. Learning to walk in virtual reality. *ACM Transactions on Applied Perception* 10, 2: 1–17. http://doi.org/10.1145/2465780.2465785
41. Joey Scarr, Andy Cockburn, Carl Gutwin, et al. 2014. The usability of CommandMaps in realistic tasks. *Proceedings of the 32nd annual ACM conference on Human factors in computing systems - CHI '14*, ACM Press, 2241–2250. http://doi.org/10.1145/2556288.2556976
42. Joey Scarr, Andy Cockburn, Carl Gutwin, and Andrea Bunt. 2012. Improving command selection with CommandMaps. *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12*, ACM Press, 257–266. http://doi.org/10.1145/2207676.2207713
43. Joey Scarr, Andy Cockburn, Carl Gutwin, and Philip Quinn. 2011. Dips and ceilings: understanding and supporting transitions to expertise in user interfaces. *Proceedings of the 2011 annual conference on Human factors in computing systems - CHI '11*, ACM Press, 2741–2750. http://doi.org/10.1145/1978942.1979348
44. Katherine Schramm, Carl Gutwin, and Andy Cockburn. 2016. Supporting Transitions to Expertise in Hidden Toolbars. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, ACM Press, 4687–4698. http://doi.org/10.1145/2858036.2858412
45. Molly E. Sorrows and Stephen C. Hirtle. 1999. The Nature of Landmarks for Real and Electronic Spaces. In *Spatial Information Theory. Cognitive and Computational Foundations of Geographic Information Science*. Springer Berlin Heidelberg, 37–50. http://doi.org/10.1007/3-540-48384-5_3
46. Perry W. Thorndyke and Sarah E. Goldin. 1983. Spatial Learning and Reasoning Skill. In *Spatial Orientation*. Springer US, Boston, MA, MA, 195–217. http://doi.org/10.1007/978-1-4615-9325-6_9
47. Md. Sami Uddin. 2016. Improving Multi-Touch Interactions Using Hands as Landmarks. Retrieved from http://hci.usask.ca/publications/view.php?id=401
48. Md. Sami Uddin and Carl Gutwin. 2016. Rapid Command Selection on Multi-Touch Tablets with Single-Handed HandMark Menus. *Proceedings of the SIGCHI Conference on Interactive Surfaces and Spaces - ISS '16*, ACM Press, 205–214. http://doi.org/10.1145/2992154.2992172
49. Md. Sami Uddin, Carl Gutwin, and Benjamin Lafreniere. 2016. HandMark Menus: Rapid Command Selection and Large Command Sets on Multi-Touch Displays. *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems - CHI '16*, ACM Press, 5836–5848. http://doi.org/10.1145/2858036.2858211
50. Norman G. Vinson. 1999. Design guidelines for landmarks to support navigation in virtual environments. *Proceedings of the SIGCHI conference on Human factors in computing systems the CHI is the limit - CHI '99*, ACM Press, 278–285. http://doi.org/10.1145/302979.303062
51. Frances A. Yates. 1966. *The art of memory*. University of Chicago Press, Chicago.
52. Bin Yu, Wei-Ying Ma, Klara Nahrstedt, and Hong-Jiang Zhang. 2003. Video summarization based on user log enhanced link analysis. *Proceedings of the eleventh ACM international conference on Multimedia - MULTIMEDIA '03*, ACM Press, 382–391. http://doi.org/10.1145/957013.957095
53. Shumin Zhai and Per-Ola Kristensson. 2003. Shorthand writing on stylus keyboard. *Proceedings of the conference on Human factors in computing systems - CHI '03*, ACM Press, 97–104. http://doi.org/10.1145/642611.642630
54. Method of loci. Retrieved September 19, 2016 from http://en.wikipedia.org/wiki/Method_of_loci
55. Downloads | Android Developers. Retrieved September 19, 2016 from http://developer.android.com/design/downloads/index.html