



HAL
open science

Effects of Locomotion and Visual Overview on Spatial Memory when Interacting with Wall Displays

Yvonne Jansen, Jonas Schjerlund, Kasper Hornbæk

► **To cite this version:**

Yvonne Jansen, Jonas Schjerlund, Kasper Hornbæk. Effects of Locomotion and Visual Overview on Spatial Memory when Interacting with Wall Displays. CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 2019, Glasgow, United Kingdom. 10.1145/3290605.3300521 . hal-02077052

HAL Id: hal-02077052

<https://hal.sorbonne-universite.fr/hal-02077052>

Submitted on 22 Mar 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Effects of Locomotion and Visual Overview on Spatial Memory when Interacting with Wall Displays

Yvonne Jansen
CNRS – Sorbonne Université
Paris, France
yvonne.jansen@sorbonne-universite.fr

Jonas Schjerlund
University of Copenhagen
Copenhagen, Denmark
schjerlund@di.ku.dk

Kasper Hornbæk
University of Copenhagen
Copenhagen, Denmark
kash@di.ku.dk

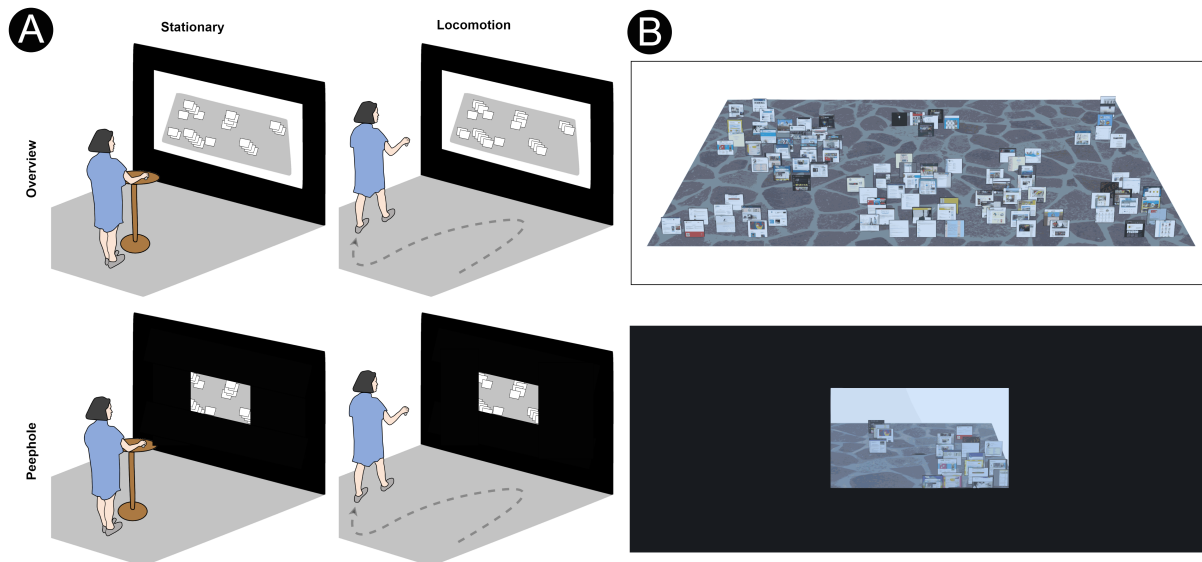


Figure 1: We adapt Data Mountain [29] to study spatial memory on wall displays. Panel A shows the four conditions in our experiment, which vary movement (stationary vs locomotion) and view (full overview vs partial peephole view). Panel B shows screenshots of two participants' layouts of web-pages in the overview (top) and peephole condition (bottom).

ABSTRACT

Wall displays support people in interacting with large information spaces in two ways: On the one hand, the physical space in front of such displays enables them to navigate information spaces physically. On the other hand, the visual overview of the information space on the display may promote the formation of spatial memory; from studies of

desktop computers we know this can boost performance. However, it remains unclear how the benefits of locomotion and overviews relate and whether one is more important than the other. We study this question through a wall display adaptation of the classic Data Mountain system to separate the effects of locomotion and visual overview. Our findings suggest that overview improves recall and that the combination of overview and locomotion outperforms all other combinations of factors.

CHI 2019, May 4–9, 2019, Glasgow, Scotland UK

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019)*, May 4–9, 2019, Glasgow, Scotland, UK

<https://doi.org/10.1145/3290605.3300521>

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI.**

KEYWORDS

Spatial memory, wall display, overview, locomotion.

ACM Reference Format:

Yvonne Jansen, Jonas Schjerlund, and Kasper Hornbæk. 2019. Effects of Locomotion and Visual Overview on Spatial Memory when Interacting with Wall Displays. In *CHI Conference on Human Factors in Computing Systems Proceedings (CHI 2019), May 4–9, 2019, Glasgow, Scotland Uk*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3290605.3300521>

1 INTRODUCTION

Spatial memory is essential to human-computer interaction, for instance, in supporting the retrieval of menu items, navigating virtual environments with confidence, or locating icons on the desktop [32]. The literature is full of examples of *desktop user interfaces* that tap spatial memory. For example, CommandMaps [33] uses the principle of spatial consistency to map the menu ribbons of Microsoft Word to the entire display space, ensuring that menu items are in the same place across invocations of CommandMaps. Using this principle gives a substantial performance gain over unmodified menu ribbons. Another successful idea is to support users in organizing objects spatially to facilitate recall; Data Mountain [29] is a classic example. In one study showing its benefits, participants were free to place 100 thumbnails of webpages on a reclined 2.5D plane; even with few cues participants were faster at recall than in a control condition.

It is less clear how spatial memory is involved in interaction with *wall displays* and how findings similar to those just cited transfer to such displays. With larger displays, many factors change. First, the increase in display size may change interaction. On wall displays, viewing angles differ even if the user is stationary and the relation between input devices and display may be hampered by having to clutch or by the decreased accuracy of absolute input. Second, in contrast to desktop displays people move when using wall displays, both in parallel to the display and away from/towards the display [e.g., 12]. Thereby, the requirements for spatial updating increase [34] and this may hamper performance and recall. Third, the scale of wall-display navigation falls between large-extent VR worlds [30] and studies of the spatial memory on desktop displays [35]; as scale is related to different kinds of spatial memory, it is not clear which findings apply to interaction with wall displays.

Together, these factors leave it unclear if wall displays may similarly tap spatial memory as do desktop displays. Specifically, the relative contribution of the spatial stability used in many desktop systems and of physical movement remains unclear. Possibly, those factors might even counteract each other. For instance, the original studies of Data Mountain used a 17-inch display. With such a display there is no need to move, the content on the display is relatively stable, and the angle relative to the display fixed. All of these factors change on a wall display.

We report on an experiment whose key idea is to tease apart the effects of having a visual overview of the information space (i.e., content on the wall display) from the effects of moving about in front of the display and in particular of being in a specific location when interacting with specific items on the display. We adapt the design of Data Mountain to wall displays and ask how the findings of Data Mountain generalize to larger displays. With that adaptation we contribute empirical data from an 80-participant study which separates the effects of locomotion and spatial overview. This study helps advance our understanding of spatial memory on wall displays and adapt guidelines for desktop displays to that setting.

2 BACKGROUND

Spatial memory has been known useful since antiquity [37] and been a focus of much research in modern Psychology. In a classic demonstration of the effect, Rothkopf studied incidental memory for location of information in text [6] and showed that people were able to recall with higher than chance accuracy where on a page they had read pieces of information. Research in spatial cognition has since solidified and grown [e.g., 34]. This research and the promise of spatial memory have inspired much work in HCI [e.g., 2, 22, 24, 39], much of it is reviewed by Scar and colleagues [32]. Next, we focus on the relation between spatially-stable overviews and spatial memory and on locomotion and its relation to spatial memory. We then discuss what overviews and locomotion mean for supporting spatial memory on wall displays.

Effects of Overviews

An overview is a spatially-stable visual representation of an entire information space. Such overviews have been created for image browsing [27], electronic maps [10], documents [24], and commands and icons [33]. Scarr and colleagues [32] recommended that “offering users a spatially-stable visual overview of an information space can greatly improve their ability to navigate to known locations, as well as reducing visual search time” (p.39). Many studies confirm the usefulness of overviews. For instance, O’Hara and colleagues [24] created an electronic document reader that provided a spatially stable view of the document page together with a detail view of individual sentences. In an empirical comparison of this interface to a traditional scrolling view of the document, they found an improved incidental memory of location.

As mentioned, the system Data Mountain [29] is another classic example of a system with a spatially-stable visual overview. Data Mountain shows thumbnails to web-pages on a reclined 2.5D plane. A user study shows that this organization of thumbnails is beneficial over a standard hierarchical list that does not use an overview. Several follow-up papers

to the original 1998 paper have shown this finding to be robust: Even in a follow-up several months after the main study, participants achieved similar recall rates as in the original study after a few blocks of trials [7].

The spatial stability may be provided in other ways than visually. Some studies have shown that the body may work as a spatially-stable overview to which information may be related [2]. Others have employed the ancient memory technique Method of Loci [37] to assign commands to physical objects in the environments [26]. Method of Loci has also recently been used in virtual reality [16]. All of these studies show substantial improvements in spatial memory from the spatially-stable overviews.

Effects of Locomotion

The effects of body movements, particularly locomotion, on spatial cognition have been studied extensively in the context of virtual reality environments [e.g., 30, 31]. Results consistently indicate that body-based movement positively affects people's ability to navigate in virtual worlds and to internalize the spatial organization of such worlds.

Large display environments also afford physical navigation as a means to explore an information space. Some positive effects of locomotion in front of these displays have been reported for a geospatial data exploration task [1], a collaborative game [19], and a sorting task [12]. Nonetheless, large display environments differ from virtual reality environments as the display is generally exocentric (i.e., fixed relative to the position and orientation of any single user), while locomotion can be mapped to either egocentric interactions (e.g., body lenses) or exocentric physical interactions (e.g., physical "zooming" which only affects a user's position relative to the display but not the display content itself). Therefore it is not clear that the benefits identified for virtual environments apply to locomotion as part of interaction on wall displays.

Peephole pointing refers to the idea that the physical movement of a display controls the navigation in a large, virtual information space [8, 38]. This typically happens without the need for a visual overview. Thus users experience the navigation centered around the point of view. Earlier work has shown performance benefits of peephole pointing over other ways of navigation, in particular, panning where the display is fixed [20, 38]. Of interest here, however, is mainly the effects of peephole pointing on spatial memory. Kaufmann and Ahlström [15] compared a projector phone with a peephole interface to a touch interface. They found no difference in performance but found substantial differences in spatial memory. Participants using the peephole interface had 41% more accurate immediate recall of seven locations that they had navigated. Device orientation does not seem to matter

for these benefits: Müller and colleagues [22] found no differences in performance or spatial memory when participants held a tablet peephole either vertically or horizontally.

Spatial memory has also been linked to kinesthetic memory, typically so that input techniques that directly map kinesthetic cues to content (e.g., touch displays) perform better than input techniques without such mappings (e.g., mouse). Tan and colleagues [36] showed that participants had 19% higher incidental recall of a set of 30 objects replaced after a 10-minute distractor task. However, this benefit of spatial memory relies on spatial stability. Jetter and colleagues [14] showed that spatial memory is supported only in panning interfaces but not in interfaces where the user can also zoom. After navigating a series of eight items, participants had to reproduce one such series. They made fewer mistakes in doing so but only using a non-panning interface.

Liu and colleagues [18] compared performance on a classification task on a wall display requiring physical zooming with a desktop interface using virtual zooming and found an interaction effect with task difficulty. The more difficult the task was, the more participants would benefit from wall display interaction.

Spatial Memory and Wall Displays

A 2013 review of spatial memory [32] did not mention wall displays and the role of spatial memory in interaction with large displays.

A few papers since then have linked spatial memory on wall displays with peephole pointing [22, 28]. Rädle and colleagues [28] investigated spatial memory in peephole pointing for wall displays using egocentric navigation with a tablet or using multi-touch while seated. They found a large decrease in path length and task time in the egocentric condition. Peephole pointing, however, makes the handheld display the focus of attention, rather than the wall display. Thus, the uses of wall displays in the literature that do not use a handheld device (for instance, because they use touch [13] or mid-air pointing [23]) are not informed by findings on peephole pointing. It is not well understood whether earlier findings apply to wall displays and, if so, how designs could tap spatial memory for wall displays.

Zagermann and colleagues [39] compared spatial memory with 11- and 55-inch displays. The experimental task was to spread out 18 letters on a canvas, first placing all the letters and then locating the letters again. In this study, the main focus was to show an impact of input modality (say, touch or pad); the findings about display size were not straightforward. Thus, the findings about spatial memory and wall displays are inconclusive.

3 EXPERIMENT RATIONALE

Research Questions

We are interested in understanding the effects of visual overview and locomotion on spatial memory. Specifically, we are interested in quantifying the effects individually as well as in symbiosis. Given the related work just reviewed, it is not clear if these effects add up or if they could even counteract each other, for instance, because physical movement could interfere with having an overview. We have participants perform a recall task of a self-organized information space on a wall display using different interaction techniques that fully crosses the availability of visual overview and locomotion. We aim to answer these questions:

- RQ1: If a wall display shows a visual overview, does it help people remember contents similar to how it supports them on desktop displays?
- RQ2: If people benefit from a visual overview of their information space, can locomotion help them further to remember where they placed items?
- RQ3: If no visual overview is provided, can locomotion help people maintain an internal map of the information space which allows them to retrieve items as reliably as when visual overview is offered?

Adapting Data Mountain to a Wall Display

Our experiment application is designed after the classic Data Mountain application [29] (see screenshot in Figure 1). Moving Data Mountain to a wall display allows two things. First, we can create a visual overview of the information space by scaling up the original Data Mountain idea to a wall display. Additionally, we can offer a partial view of the information space where only a small portion of an area with its immediate neighbors is visible at any point in time.

Second, the inclined plane of our Data Mountain implementation maps naturally to the floor in front of the wall display, as coordinates on the inclined plane of the Data Mountain interface map to a physical location in front of it. Thus, a participant has to move to the physical location corresponding to the virtual plane coordinate occupied by an item to interact with it (see Figure 2). Alternatively, they are stationary and use an input device to move items.

In this way, our adaptation of Data Mountain provides a way to manipulate visual overview and locomotion fully crossed. As in the Data Mountain, we use relative input with a mouse; we decided against touch input or freehand input because both might influence memory unpredictably.

4 EXPERIMENT

The overall goal is to examine how visual overview and locomotion influence retrieval performance and the strategies people use to organize information and thus infer the role

each play in forming and using spatial memory of items in an information space.

Design

Participants use one of four conditions (between-participants design), varying the view and movement availability of visual overview and locomotion in a fully crossed design. We call the availability of visual overview *overview*, the absence *peephole*; the availability of locomotion is *locomotion*, the absence *stationary*. The conditions are as follows (see Figure 1):

- *Overview+locomotion*: Participants physically navigate a visually stable overview of the information space on a wall display. This condition resembles a typical scenario of using wall displays for working with large information spaces [13, 23], except that participants have to move.
- *Overview+stationary* (i.e., overview only): Participants move a cursor in a visually stable overview using a mouse while standing at a fixed location in front of the wall display, thus preventing possible benefits of locomotion (e.g., relating item locations to locations on the floor). This condition resembles the original Data Mountain. [29], except that the much larger display may influence spatial cognition [35].
- *Peephole+locomotion* (i.e., locomotion only): Participants move a peephole view through locomotion, which may help orient in the information space. This condition resembles the egocentric condition of Rädle and colleagues [28], except participants in that study focused their attention on a tablet and not a wall display.
- *Peephole+stationary* (a baseline with no overview and no locomotion): Participants move a peephole view




Figure 2: Floor-to-display mapping. In the locomotion conditions, users' position on the floor determine what they may interact with.

using a mouse while at a fixed location. This condition resembles a typical small-view interface, although it does not support zooming to suppress the option of seeing an overview of the entire information space.

Apparatus

The study runs on a Windows 10-based PC with an NVidia GeForce 1080 GPU, an Intel i7-8700K CPU @ 3.7GHz, and 16GB of DDR4 RAM. It drives a Dell S718QL 4K laser projector that projects onto a 2.8m×1.172m 2.39:1 anamorphic wall display, downscaled to a 3500×1465 resolution to fit the display's aspect ratio. The software that drives the display is implemented in Unity 5.6.2f1.

We use a NaturalPoint OptiTrack motion capture system (www.naturalpoint.com/optitrack/) with 24 ceiling-mounted infrared cameras to determine the cursor's position in the locomotion-based conditions. More specifically, we track four infrared markers mounted on the Gyratation Air Mouse Go Plus wireless mouse that we use for the study; this mouse was also used in the *stationary* conditions.

Interface

The interface allows interacting with an information space that resembles Data Mountain [29]: It provides a perspective view of a plane (tilted at 40 degrees on the X-axis) on which items (web bookmarks in Data Mountain, screenshots of websites in the present study) can be placed.

There are two versions of the interface corresponding to the manipulation of *overview*: a stable visual overview of the entire information space, in which a cursor can be moved around (Figure 1, left) and a peephole where a view port can be panned to show only part of the information space at a time (Figure 1, right). In the conditions with a visual overview (i.e., *overview+locomotion* and *overview+stationary*), the entire information space fits within the display and remains in a fixed centered position. In the peephole conditions, (*peephole+locomotion* and *peephole+stationary*), a 16:9 window of 1050×590 pixels (out of a total resolution of 3500×1050 pixels) is shown in the center of the display. This window shows the part of the information space around the cursor's location. When the user moves the mouse in the plane, the view is panned so that the cursor remains in a fixed position within the view.

In all conditions, users can interact with content in three ways: (i) add a new item, (ii) move an item, and (iii) bring up a detail view to inspect an item:

- Users can add new items at the cursor's location by right-clicking while the cursor is in an empty location in the information space. Adding a new item automatically brings up the detail view of the item, which can be dismissed with a right-click.

- When the cursor is hovering over an item, the item is highlighted with a white outline. The user can select a highlighted item by left-clicking, and then drag the item by moving the cursor while holding down the mouse button. The movement of the item is continuous and constrained to the surface of the plane. Data Mountain allowed to push items with the one being currently dragged in order to avoid occlusion. However, this results in unstable information spaces leaving items at potentially different locations than where a participants placed it (thus possibly interfering with their spatial memory). We thus decided to instead set a threshold distance at which the item being moved towards another item stops following the cursor at one side of that item and then jumps directly to the other side once the cursor has moved beyond the threshold distance.
- Users can bring up a zoomed detail view of a highlighted item for closer inspection (similar to moving items forward to a preferred viewing position in Data Mountain) by right-clicking; right-clicking again dismisses the detail view.

Control of the movement of the on-screen cursor differs between conditions. Cursor movements in the *locomotion* conditions (*overview+locomotion* and *peephole+locomotion*) are controlled by directly mapping the user's position in front of the display to the information plane (see Figure 2). In the *stationary* conditions (*overview+stationary* and *peephole+stationary*), the cursor maps to the mouse movement. In contrast to Data Mountain, the cursor is represented as a black circle on the plane. We kept the inclined plane from Data Mountain to help participants understand the floor-to-display mapping.

Stimuli and Tasks

The stimuli used in the experiment were 100 screenshots of web pages with legible text, randomly selected from the material provided by Jahanian et al. [11]. We chose these stimuli because they were high in resolution and recent, yet similar to the stimuli used in the Data Mountain studies.

These stimuli were used in a placement phase, where participants had to organize 100 screenshots on the inclined plane, and in a retrieval phase, where participants saw a fixed random subset of the placed items one-by-one and had to indicate their location; this is similar to the tasks in Data Mountain [29].

Measures

The Data Mountain study only measured how many items were correctly identified but allowed at the same time several attempts before a timeout. Since repeated retrieval attempts

can serve to consolidate memory upon feedback, we opted instead for a more continuous measure, the Euclidean distance, that is, the straight-line distance, from the indicated item to the correct item. As additional measures, we capture how many items are closer to the correct item than the selected one and how many items were correctly selected.

We collected measures of participants' perception of their effort after placing and after retrieving items using NASA's Task Load Index (TLX, [9]).

We also collect data to capture potential individual behavior differences in placing and retrieving items including (i) the total time spent placing items, (ii) how long participants looked at the detail view of an item before shrinking it down, (iii) the space used to place items, measured as world-space X- and Z-axis coordinates in the study application, (iv) the total duration of the retrieval phase.

Finally, we collect individual measures of the visual working memory capacity measured as described by Pailian and colleagues ([25], experiment 4A). We used the software platform PEBL2 [21] and a custom experiment script. The rationale is that visual working memory capacity is likely a predictor of variation in individual performance.

Participants

Eighty participants (36 female) were recruited through advertisements in local Facebook groups as well as through a mailing list of people interested in signing up for experiments. Participants had to be at least 18 years old and were instructed to complete an online test for color blindness before signing up. This test was re-administered before the experiment to ensure full color vision. Participants were paid the equivalent of about €20 for agreeing to participate.

Procedure

Participants were randomly assigned to one of the four conditions in order of arrival. The experiment had three main phases: In the *placement phase*, participants were asked to organize 100 items. Participants could move and zoom existing items and add new items one at a time. Participants were only made aware of the later need to recall the items with the formulation that "there will be a later part in the experiment where you will reuse your organization for other tasks, including re-finding specific items" (the complete wording can be found in the experimenter instructions <https://osf.io/au7tf>). This phase lasted on average 28 minutes.

In the *filler phase*, participants were seated in front of a separate 24-inch display and for 15 minutes completed the visual working memory capacity test.

In the *retrieval phase*, participants were shown the organization of the items they created in the storage phase, but with the contents of the items greyed out. Participants had to recall the location of 25 fixed random items, given as one

target cue at a time. They were instructed as follows: "Try to remember where you placed it and left click on the corresponding placeholder. If you don't remember exactly where you placed it, click on a placeholder in the vicinity of the area you think is most likely." The target cue was 300×300 pixels and displayed at all times in the top-left corner of the wall display. The cue could be zoomed using an identical functionality as seen in the storage phase. Participants received feedback whether their choice was correct.

Then the next item would be shown. This phase lasted on average 4.3 minutes.

Preregistration and Statistical Analysis

Cockburn et al. [4] recently reviewed the many benefits of preregistering one's hypotheses, study plans, and data analysis intents. The purpose of such a preregistration is to increase the transparency of research intents and to enable the reader to verify that the research questions and hypotheses presented in an academic paper were not adapted to the data after analysis. We preregistered our data analysis approach using the persistent data storage of the Open Science Framework (OSF) <https://osf.io/pkqbj> and also make additional material such as experimenter instructions and our data available via the project page <https://osf.io/2rxfn/>.

Summarizing our analysis plan, we compute bootstrapped 95% confidence intervals for all measures, per condition and factor. We exclude participants if performance on any of the individual-difference variables differs by more than three standard deviations which, assuming an underlying normal distribution, only excludes the most extreme 0.2%. We excluded nine participants based on these criteria: two due to time taken to complete the placement phase, six due to how much of the available space they used, and one due to time taken to complete the retrieval phase.

5 RESULTS

We present our results according to the planned analysis starting with the mean measure, *Euclidean distance*, followed by the auxiliary measures, number of *items correctly selected*, and number of *items closer to correct target*. Then we report on five control measures which served to identify outlier participants. The figures indicate estimates for the means of the reported measures with 95% bootstrapped confidence intervals, i.e. the error bars indicate the range of values within which the true means can be expected with 95% confidence. The respective sub-figures to the right indicate either pairwise differences of the two measures to their left or of differences between factors. These differences indicate estimates of effect sizes in the original unit of the measures or factors. The further an error bar is from the line indicating "0", the stronger the data provide evidence for a clear effect of a factor.

Main Measure

Figure 3 shows the results of our main measure, the Euclidean distance to the correct target for our two factors, view and movement. We expected that a visual overview would help participants to consolidate their memory of the layout and that those experiencing the overview level would outperform those in the peephole level. Similarly, we expected that movement would have a positive effect on spatial memory and that participants using locomotion would outperform those being stationary. As Figure 3 indicates, the data provides evidence for an effect of view. Concerning movement the evidence is insufficient to determine whether it has a clear effect. We can only say that our results point in the expected direction.




Figure 3: Factor-wise contrasts for the main measure, Euclidean distance to correct target together with differences between factor levels. Error bars represent 95% bootstrapped confidence intervals.

If we now look at the data at a factor combination level, we find that one combination, *overview+locomotion* outperformed all other factor combinations (see Figure 4) while the other combinations performed somewhat similarly.




Figure 4: The main measure, Euclidean distance to correct target, per condition together with pairwise differences to the overview+locomotion condition. Error bars represent 95% bootstrapped confidence intervals.

Auxiliary Measures

As indicated in the preregistration document, we planned to consolidate our findings by also measuring how many items were correctly selected, and how many items were closer to the correct target. The rationale behind these extra measures is that we observed in a pilot that a possible strategy is to place items closely together which can lead to misleadingly small Euclidean distances.

For the number of *items correctly selected* measure (see Figure 5), we find that the data indicate that a possible effect is more likely in the predicted direction (top). However, the differences are too small and the estimation too imprecise to definitely determine that either factor has an effect on correctly selected items.

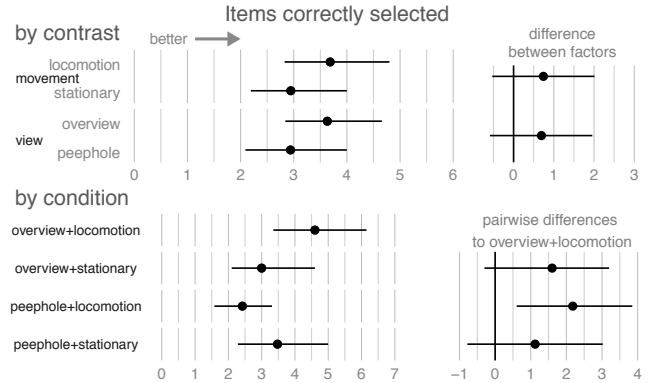


Figure 5: Items correctly selected as factor-wise contrasts (top) with differences between factor levels and per condition (bottom) with pairwise differences. Error bars represent 95% bootstrapped confidence intervals.

However, looking at the data per condition (bottom) we find again, as with the *distance* measure, that *overview+locomotion* performs better than other factor combination, although the data here are noisier, and the evidence is weaker for two of the pairwise differences.

Results for the second auxiliary measure, number of *items closer to correct target*, are shown in Figure 6. Similar to the *distance* measure, the contrasts indicate an effect for view but not for movement (top). The comparison per condition (bottom) indicates again that *overview+locomotion* performs better than the other factor combinations.

In summary, across the three dependent variables, we find consistently an effect of view indicating that *overview* is the most relevant feature to improve spatial memory. In addition, we observed that locomotion seems to play a certain role as repeatedly the *overview+locomotion* condition outperformed the other factor combinations. However, just looking at the contrasts, the effect size of movement is too small to be estimated accurately given the statistical power of our experiment.

Standardized effect sizes. Additionally to the planned analysis, we include in Figure 7 standardized effect sizes for the main and the two auxiliary measures, respectively for the two factors and the interaction term between them. Additionally, we computed unit weights for the three measures (z-scores) which allow us to create a weighted composite score from the sum of the unit weights [3]. This composite




Figure 6: Items closer to correct target as factor-wise contrasts (top) with differences between factor levels and per condition (bottom) with pairwise differences. Error bars represent 95% bootstrapped confidence intervals.

score (shown on the bottom) indicates a reliable effect of view across the three measures and still large uncertainty as to the likely effect size of movement. A relevant interaction effect seems rather unlikely as most of the confidence interval falls between $d = -0.2$ and $d = 0.2$.

Control Measures

To investigate whether the effect of view and more specifically the apparent advantage for *overview+locomotion* can be attributed to our manipulation or whether it might be due to some idiosyncratic behavior emerging in this factor combination, we captured five control measures to detect individual differences used to exclude individual participants and to detect whether the random assignment happened to




Figure 7: Standardized effect sizes for both factors and their interaction for the main and auxiliary measures as well as a weighted composite score. Error bars represent 95% bootstrapped confidence intervals.

result in considerable group differences. We captured the following variables for this purpose: time spent during the two phases, the use of available space, and the time spent inspecting the targets in detail view. Here we only present textual summaries for these measures, additional figures are available on the OSF platform at <https://bit.ly/2NvzCzK>.

Time spent in the placement phase. Participants in the *overview+locomotion* condition spent on average slightly more time placing items (30.5 min [27.1, 33.8]) than participants in the *overview+stationary* (25.8 [21.0, 31.0]), *peephole+locomotion* (26.2 [19.2, 33.2]) or *peephole+stationary* condition (25.6 [22.9, 28.5]).

Time spent inspecting targets in detail view. Concerning this measure, we find that participants in the *peephole+locomotion* condition spent a bit less time (3.5 sec [2.7-4.5]) looking at items in detail than participants in one of the overview conditions (*overview+locomotion* 5.3 [4.4, 6.1], *overview+stationary* 5.3 [4.3, 6.6], *peephole+stationary* 4.5 [3.5, 5.4]).

Time spent in the retrieval phase. The time spent to retrieve items was comparable between conditions. The point estimates are all in the range 4.1-4.3 min [3.4-3.7, 4.8-5.0].

Use of space. Concerning the use of space, we find no notable differences between conditions on the level of the bounding rectangle of all placed items. Most participants, across conditions, used most of the available space: they placed at least one item close to each edge of the inclined plane. Still, this measure led to the most exclusions of participants: six people were excluded because they used a lot less space than most people. Figure 8 gives an overview of how participants, including outliers, laid out their information spaces (note that participants 51, 55, and 56 were excluded for other reasons). It is difficult to perform analysis for these layouts, but let us make two observations. First, the organization of content seems to happen mainly in terms of rows (e.g., participants 3, 17, 32), columns (e.g., participants 1, 2, 7, 28), and clusters (e.g., 4, 6, 31, 53). Second, there is large variations in the screen real estate used for content; compare, for instance, Participant 7 to Participant 69. It remains open work to relate these organizations to input conditions and to retrieval strategies; one hypothesis could be that the Overview encourages more use of clusters whereas Peephole encourages in particular column organization.

Visual working memory capacity. We measured participants' visual working memory capacity using the flickering technique described by Pailian and Halberda [25]. The test was used as a filler task between the two phases (placement and retrieval). Surprisingly, we find that participants in the *peephole+stationary* condition appear to have a much lower capacity than participants in other conditions (see Figure 9).




Figure 8: All layouts created by participants. Those shown in grey belong to participants excluded from the analysis. Numbers indicate participant ids. Layouts 26 and 53 are used in Figure 1-B.




Figure 9: Visual working memory capacity for five items. Error bars represent 95% bootstrapped confidence intervals.

This is quite unlikely to happen with random assignment of participants to conditions, and we believe it is more likely an indicator that this condition was more demanding and that participants were less attentive while doing the task.

NASA TLX. The collection of these data were not part of the preregistration but uses the same analysis methods. The summary of the results is shown in Figure 10 both for the subscales and the weighted overall score. The data show no surprises except that the worse scores for the visual working memory in the *peephole+stationary* condition were apparently not due to *conscious* extra effort in this condition: participants did not report higher frustration, effort, or mental demand in this condition than in others.

6 DISCUSSION

We were interested in the individual and joint effects of locomotion and visual overview for wall displays. Our main result is that the visual overview improves recall, measured as the Euclidean distance from the indicated position to the target item across 25 targets. Moreover, the combination of visual overview and locomotion when placing and recalling items had higher recall success over all other conditions. We did not find conclusive evidence for other differences across conditions for the recall measure. Participants spend slightly more time placing items in the *overview+locomotion* condition but the spatial characteristics of their layouts and the time used to retrieve items were similar across conditions.

The Influence of View and Locomotion

We find that having an overview improves recall over the peephole condition: View has a medium-size effect of Cohen's $d = .46$, [0.02, 0.94] 95% CI. This improvement is the key finding of our paper. It confirms the importance in general of spatial stability [32], and it confirms that a visual overview like the Data Mountain [29] benefits spatial memory also on wall displays. It seems that even if an overview is




Figure 10: Results for the NASA TLX questionnaire for both phases. The raw scores have been mapped to a 0-100 scale to aid comparisons with the weighted overall score. Error bars represent 95% bootstrapped confidence intervals.

substantially larger than in related work (such as the 55-inch display studied by Zagermann and colleagues [39]), spatial memory is still supported. Thus, our first research question (RQ1) is answered affirmatively.

For movement as a factor, our data provide insufficient evidence to determine a clear effect. The data suggest that a possible effect might go in the predicted direction – locomotion having a positive effect – but that the effect is at best too small to be detectable with the statistical power of our experiment. A much larger sample size is required to provide a more definitive answer.

A curious finding is that the combination of locomotion and overview outperforms all other factor combinations. On average, participants are 1.5 units (or 30%) closer, in Euclidean distance, to the correct target with *overview+locomotion*. This finding provides evidence in favour of our second research question (RQ2) and suggests an important synergy between overview and locomotion. It may also be interpreted as confirming earlier work, which found that locomotion is helpful [1, 12]. It might be a consequence of our design decision that physical movement is not only encouraged (as in [1]) but required to select and move items. However, our statistical power was insufficient to accurately estimate the effect size of the movement factor.

One mechanism behind the high recall with the *overview+locomotion* condition might be that the layout phase is more thorough, as suggested by the approximately five extra minutes spent in the *overview+locomotion* condition. It could be

that locomotion simply slows down placement and that slowing down is particularly useful when simultaneously seeing the overview. Another possible mechanism is that dual spatial encoding of items occurs by being at a particular location in front of the display when in the *locomotion* conditions but at the same time seeing a particular stable, spatial location on the wall display (in the *overview* conditions). Again, the benefit only arises with the combination. We have no data to confirm this speculation about mechanisms. However, the worry—expressed earlier in this paper—that those two spatial encodings might interfere with each other, is not warranted.

Stationary Peephole Worked Better than Expected

A third finding of this paper is that the absence of an overview in combination with the absence of movement performed surprisingly well. We expected this combination to perform the worst of all conditions. It was intended as a baseline in comparison to which all other conditions would be improvements. This expectation was not met.

One explanation might be that peephole displays just work well, as suggested by related work [e.g., 15, 20]. Zagermann et al. [39] found a mixed effect of display size, which might suggest difficulties in effectively using the *overview* conditions. Our findings are similar to a study of map overviews [10] which found that a detail-only zoomable interface worked well compared to an interface with an overview.

Of course, the finding discussed here is an absence of evidence for a difference that we expected; it is therefore

somewhat weak. Thus, our third research question (RQ3) remains open.

Limitations

Our study has a number of limitations and differences to related work. First, our measure of recall differ from that used in the Data Mountain studies [e.g., 29] and can thus not be directly compared to the previous results.

Second, locomotion may be implemented in many ways. We use a simple mapping from the floor to the display (as shown on Figure 2). Given the importance of locomotion established in earlier work [e.g., 1, 13], it is surprising that more guidelines to help our design choices were not available. This leaves plenty of room for future work to develop other uses of locomotion in front of wall displays.

Third, our study does not evaluate whether or not the design choices behind Data Mountain are the best way of implementing spatial memory support in general. In particular, it is controversial what the inclined plane and 2.5D components of Data Mountain add to memory performance [5, 6]. In particular, it remains an open question what the effect of the 2.5D components of our implementation have on performance and whether those components could be designed more effectively.

Fourth, in comparison to other memory studies in HCI [e.g., 1, 13, 28], and to work on large displays more generally [e.g., 17], we have reported comparatively little on interaction. One key open question about interaction, that we have not seen discussed in earlier work, concerns the influence on recall of the actual spatial organization that participants use. Figure 8 suggests several patterns in organization that future work should look into, trying to establish their influence on performance and memory.

7 CONCLUSION

Earlier work has shown that spatially stable overviews and locomotion contribute to improved spatial memory; we have investigated if this is also true for wall displays and identify an effect of overview and a combined effect of overview and locomotion. Our findings suggest that overview improves recall and that the combination of overview and locomotion outperforms all other factor combinations.

8 ACKNOWLEDGEMENTS

We thank Mikkel Rønne Jakobsen for his contributions during early phases of this research. This work was supported by the European Research Council, grant no 648785.

REFERENCES

- [1] Robert Ball, Chris North, and Doug A Bowman. 2007. Move to Improve: Promoting Physical Navigation to Increase User Performance with Large Displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*. ACM, New York, NY, USA, 191–200. <https://doi.org/10.1145/1240624.1240656>
- [2] Joanna Bergstrom-Lehtovirta, Sebastian Boring, and Kasper Hornbæk. 2017. Placing and Recalling Virtual Items on the Skin. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 1497–1507. <https://doi.org/10.1145/3025453.3026030>
- [3] Philip Bobko, Philip L Roth, and Maury A Buster. 2007. The usefulness of unit weights in creating composite scores: A literature review, application to content validity, and meta-analysis. *Organizational Research Methods* 10, 4 (2007), 689–709.
- [4] Andy Cockburn, Carl Gutwin, and Alan Dix. 2018. HARK No More: On the Preregistration of CHI Experiments. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 141, 12 pages. <https://doi.org/10.1145/3173574.3173715>
- [5] Andy Cockburn and Bruce McKenzie. 2001. 3D or Not 3D?: Evaluating the Effect of the Third Dimension in a Document Management System. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*. ACM, New York, NY, USA, 434–441. <https://doi.org/10.1145/365024.365309>
- [6] Andy Cockburn and Bruce McKenzie. 2002. Evaluating the Effectiveness of Spatial Memory in 2D and 3D Physical and Virtual Environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '02)*. ACM, New York, NY, USA, 203–210. <https://doi.org/10.1145/503376.503413>
- [7] Mary Czerwinski, Maarten Van Dantzich, George G Robertson, and Hunter G Hoffman. 1999. The Contribution of Thumbnail Image, Mouse-over Text and Spatial Location Memory to Web Page Retrieval in 3D.. In *INTERACT*. 163–170.
- [8] George W. Fitzmaurice. 1993. Situated Information Spaces and Spatially Aware Palmtop Computers. *Commun. ACM* 36, 7 (July 1993), 39–49. <https://doi.org/10.1145/159544.159566>
- [9] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183.
- [10] Kasper Hornbæk, Benjamin B Bederson, and Catherine Plaisant. 2002. Navigation patterns and usability of zoomable user interfaces with and without an overview. *ACM Transactions on Computer-Human Interaction (TOCHI)* 9, 4 (2002), 362–389.
- [11] Ali Jahanian, Shaiyan Keshvari, and Ruth Rosenholtz. 2018. Web pages: What can you see in a single fixation? *Cognitive Research: Principles and Implications* 3, 1 (09 May 2018), 14. <https://doi.org/10.1186/s41235-018-0099-2>
- [12] Mikkel R Jakobsen and Kasper Hornbæk. 2015. Is Moving Improving?: Some Effects of Locomotion in Wall-Display Interaction. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 4169–4178. <https://doi.org/10.1145/2702123.2702312>
- [13] Mikkel R Jakobsen, Yvonne Jansen, Sebastian Boring, and Kasper Hornbæk. 2015. Should I stay or should I go? Selecting between touch and mid-air gestures for large-display interaction. In *Human-Computer Interaction*. Springer, 455–473.
- [14] Hans-Christian Jetter, Svenja Leifert, Jens Gerken, Sören Schubert, and Harald Reiterer. 2012. Does (Multi-)Touch Aid Users' Spatial Memory and Navigation in 'Panning' and in 'Zooming & Panning' UIs?. In *Proceedings of the International Working Conference on Advanced Visual Interfaces (AVI '12)*. ACM, New York, NY, USA, 83–90. <https://doi.org/10.1145/2254556.2254575>
- [15] Bonifaz Kaufmann and David Ahlström. 2013. Studying Spatial Memory and Map Navigation Performance on Projector Phones with Peephole Interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 123–132. <https://doi.org/10.1145/1240624.1240656>

- Factors in Computing Systems (CHI '13)*. ACM, New York, NY, USA, 3173–3176. <https://doi.org/10.1145/2470654.2466434>
- [16] Eric Krokos, Catherine Plaisant, and Amitabh Varshney. 2018. Virtual memory palaces: immersion aids recall. *Virtual Reality* (16 May 2018). <https://doi.org/10.1007/s10055-018-0346-3>
- [17] Lars Lischke, Sven Mayer, Katrin Wolf, Niels Henze, Albrecht Schmidt, Svenja Leifert, and Harald Reiterer. 2015. Using Space: Effect of Display Size on Users' Search Performance. In *Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA '15)*. ACM, New York, NY, USA, 1845–1850. <https://doi.org/10.1145/2702613.2732845>
- [18] Can Liu, Olivier Chapuis, Michel Beaudouin-Lafon, Eric Lecolinet, and Wendy E Mackay. 2014. Effects of display size and navigation type on a classification task. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*. ACM, 4147–4156.
- [19] Sven Mayer, Lars Lischke, Jens Emil Grønbaek, Zhanna Sarsenbayeva, Jonas Vogelsang, Paweł W. Woźniak, Niels Henze, and Giulio Jacucci. 2018. Pac-Many: Movement Behavior when Playing Collaborative and Competitive Games on Large Displays. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18)*. ACM, New York, NY, USA, Article 539, 10 pages. <https://doi.org/10.1145/3173574.3174113>
- [20] Sumit Mehra, Peter Werkhoven, and Marcel Worring. 2006. Navigating on Handheld Displays: Dynamic Versus Static Peephole Navigation. *ACM Trans. Comput.-Hum. Interact.* 13, 4 (Dec. 2006), 448–457. <https://doi.org/10.1145/1188816.1188818>
- [21] Shane T Mueller and Brian J Piper. 2014. The psychology experiment building language (PEBL) and PEBL test battery. *Journal of neuroscience methods* 222 (2014), 250–259.
- [22] Jens Müller, Roman Rädle, Hans-Christian Jetter, and Harald Reiterer. 2015. An Experimental Comparison of Vertical and Horizontal Dynamic Peephole Navigation. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15)*. ACM, New York, NY, USA, 1523–1526. <https://doi.org/10.1145/2702123.2702227>
- [23] Mathieu Nancel, Julie Wagner, Emmanuel Pietriga, Olivier Chapuis, and Wendy Mackay. 2011. Mid-air pan-and-zoom on wall-sized displays. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 177–186.
- [24] Kenton O'Hara, Abigail Sellen, and Richard Bentley. 1999. Supporting Memory for Spatial Location While Reading from Small Displays. In *CHI '99 Extended Abstracts on Human Factors in Computing Systems (CHI EA '99)*. ACM, New York, NY, USA, 220–221. <https://doi.org/10.1145/632716.632853>
- [25] Hrag Pailian and Justin Halberda. 2015. The reliability and internal consistency of one-shot and flicker change detection for measuring individual differences in visual working memory capacity. *Memory & Cognition* 43, 3 (01 Apr 2015), 397–420. <https://doi.org/10.3758/s13421-014-0492-0>
- [26] 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. In *Proceedings of the 33rd annual acm conference on human factors in computing systems*. ACM, 299–308.
- [27] Catherine Plaisant, David Carr, and Ben Shneiderman. 1995. Image-browser taxonomy and guidelines for designers. *Ieee Software* 12, 2 (1995), 21–32.
- [28] Roman Rädle, Hans-Christian Jetter, Simon Butscher, and Harald Reiterer. 2013. The Effect of Egocentric Body Movements on Users' Navigation Performance and Spatial Memory in Zoomable User Interfaces. In *Proceedings of the 2013 ACM International Conference on Interactive Tabletops and Surfaces (ITS '13)*. ACM, New York, NY, USA, 23–32. <https://doi.org/10.1145/2512349.2512811>
- [29] George Robertson, Mary Czerwinski, Kevin Larson, Daniel C Robbins, David Thiel, and Maarten van Dantzich. 1998. Data Mountain: Using Spatial Memory for Document Management. In *Proceedings of the 11th Annual ACM Symposium on User Interface Software and Technology (UIST '98)*. ACM, New York, NY, USA, 153–162. <https://doi.org/10.1145/288392.288596>
- [30] Roy A Ruddle. 2013. The Effect of Translational and Rotational Body-Based Information on Navigation. In *Human Walking in Virtual Environments*. Springer New York, 99–112. https://doi.org/10.1007/978-1-4419-8432-6_5
- [31] Roy A Ruddle, Ekaterina Volkova, Betty Mohler, and Heinrich H Bühlhoff. 2011. The effect of landmark and body-based sensory information on route knowledge. *Memory & cognition* 39, 4 (May 2011), 686–699. <https://doi.org/10.3758/s13421-010-0054-z>
- [32] Joey Scarr, Andy Cockburn, Carl Gutwin, et al. 2013. Supporting and exploiting spatial memory in user interfaces. *Foundations and Trends® in Human-Computer Interaction* 6, 1 (2013), 1–84.
- [33] Joey Scarr, Andy Cockburn, Carl Gutwin, and Andrea Bunt. 2012. Improving command selection with CommandMaps. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 257–266.
- [34] Priti Shah and Akira Miyake. 2005. *The Cambridge handbook of visuospatial thinking*. Cambridge University Press.
- [35] Desney S Tan, Darren Gergle, Peter Scupelli, and Randy Pausch. 2006. Physically large displays improve performance on spatial tasks. *ACM Transactions on Computer-Human Interaction (TOCHI)* 13, 1 (2006), 71–99.
- [36] Desney S. Tan, Randy Pausch, Jeanine K. Stefanucci, and Dennis R. Proffitt. 2002. Kinesthetic Cues Aid Spatial Memory. In *CHI '02 Extended Abstracts on Human Factors in Computing Systems (CHI EA '02)*. ACM, New York, NY, USA, 806–807. <https://doi.org/10.1145/506443.506607>
- [37] Frances Amelia Yates. 2013. *Art of Memory*. Routledge.
- [38] Ka-Ping Yee. 2003. Peephole Displays: Pen Interaction on Spatially Aware Handheld Computers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '03)*. ACM, New York, NY, USA, 1–8. <https://doi.org/10.1145/642611.642613>
- [39] Johannes Zagermann, Ulrike Pfeil, Daniel Fink, Philipp von Bauer, and Harald Reiterer. 2017. Memory in Motion: The Influence of Gesture- and Touch-Based Input Modalities on Spatial Memory. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems (CHI '17)*. ACM, New York, NY, USA, 1899–1910. <https://doi.org/10.1145/3025453.3026001>