

An Evaluation of Methods for Linking 3D Views

Matthew Plumlee
Department of Computer Science
University of New Hampshire
Durham, NH 03824, USA
mdp@cs.unh.edu

Colin Ware
Data Visualization Research Lab
Center for Coastal and Ocean Mapping
University of New Hampshire
Durham, NH 03824, USA
colinw@cisunix.unh.edu

Abstract

It is common for 3D visualization systems to provide multiple points of view to a user, but there have been many solutions to the problem of linking these views so that users can understand the spatial relationships between them. Toward developing guidelines for view-linking devices, we have carried out two experiments that compare the utility of three different classes of linking devices: a directional proxy, tethers from one view to another, and a track-up map coupling. The task we apply them to is what we call the multi-perspective identification task: subjects are asked to identify an item seen in a local, forward-looking view in the context of a global, overhead view. Our results indicate that the directional proxy is the most beneficial device, and that the track-up map coupling is also beneficial. The results suggest that tethers provide little benefit. The results also suggest that when multiple local views are present, it may be beneficial to emphasize one window as being of primary interest.

Keywords

Multiple windows, interaction design, frames of reference, Multi-perspective identification task

1. Introduction

When we look at an information source, such as a picture, map, or page of text, we often cannot take all of the information in at once. We must instead shift our attention from portion to portion, identifying important objects as well as relationships between the objects. When performing a task such as guiding an underwater vehicle to a new point of interest, we often combine information from several sources (such as a sonar map of the area, a video feed from the vehicle, and estimations on the location and orientation of the vehicle). Often it is necessary to see aspects of the same geographic information space at different scales, e.g. a local view of a vehicle together with an overview giving the spatial context. If we are to create a coherent semantic picture of the relevant information in our minds, or at least be able to make reliable decisions based on such information, we need some way of establishing relationships between the information in each view.

It is possible in many instances to combine information from multiple sources into a single view, as is done with layers in GIS

systems or color and texture combinations in a number of information visualization techniques. However, there are many situations in which information cannot be integrated into a single viewpoint, for reasons of scale, perspective, information overload, or others. We focus our discussion on the latter situations, where multiple 3D windows are useful. We specifically concentrate on three methods for linking the information in each view. The first two visually express an existing relationship, while the third enforces a visual relationship that would not otherwise exist:

1. View proxy—the explicit representation of one view (or point of interest) within another. Figure 1 shows a triangle-shaped proxy indicating view direction in three dimensions, with the apex representing the eyepoint.
2. Tethers—explicit lines connecting one view (or point of interest) to its location in another. Figure 1 demonstrates a two-tether interface that connects two corners of one view window with the corresponding point of interest in the other view.
3. Orientation coupling—an implicit aid that keeps two views oriented in similar directions. The lower image in Figure 1 demonstrates a track-up coupling: an overview that keeps its “up” direction aligned with the “forward” direction of the inset view. This contrasts with the more traditional north-up overview in the upper image.

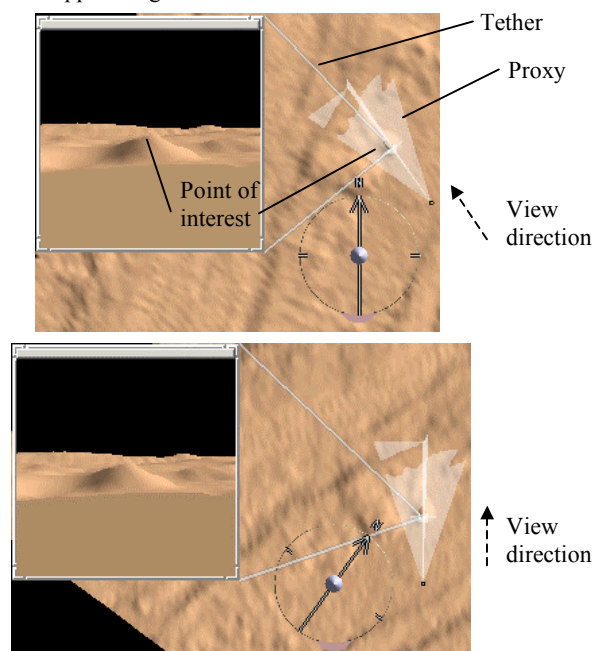


Figure 1: Examples of the three linking methods we have investigated. See text for explanation.

View proxies are integral parts of some multi-view interaction techniques, and some studies have been done already to show their utility. DragMag, by Ware and Lewis [1995], uses proxies in an overview map both to show the area covered in a zoomed view and to facilitate direct manipulation of that area by allowing the user to drag the proxy around the overview. Worlds in Miniature by Stoakley, Conway, and Pausch [1995] replicates all items seen in an egocentric view to create a miniature (exocentric) overview. The miniature overview contains a camera proxy that helps relate the two views visually, and can be moved to change the egocentric viewpoint.

Yamaashi et al. [1996] demonstrate that a “linked” proxy (one that can be used for navigating its represented window) reduces the time to perform a multiscale identification task. Their task has subjects monitor a wider-angle video display for the presence of a character in the scene, and cause a second video display to zoom in so as to identify the character. The number of user operations is reduced because moving a linked proxy in two dimensions is a single operation, while zoom, pan, and tilt operations must be composed to move a detail view around in the absence of a linked proxy. The time reduction (roughly 45%) can be accounted for by the reduced number of operations.

Tethers appear less often in computer science literature, but do appear in printed illustrations. Illustrations in magazines such as National Geographic “blow-up” a portion of an image and connect this blown-up portion to a proxy on the overview image. DragMag [Ware and Lewis 1995] uses tethers to a similar effect in connecting each mag-window to its proxy in the overview region. The Spiral Calendar of Mackinlay, Robertson, and DeLine [1994] uses semitransparent, planar tethers to connect calendars on successively longer time scales. These tethers help to point out one calendar’s location in the context of the next wider-scaled calendar. The Starlight system described by Risch et al. [1997] uses tethers in a different way, through a construct they call tie-nodes. Tie nodes connect different representations of the same object in respective views or contexts.

The use of an overview map can dramatically improve our understanding of geographical space and short circuit the laborious process of learning geography by navigating through it. Seigel and White [1975] define three kinds of spatial knowledge—declarative (landmarks), procedural (routes), and survey (maps) knowledge. Thorndyke and Hayes-Roth [1982] performed experiments to compare the utility of maps and route (“navigation”) experience in a number of situations. Their findings show that 20 minutes of studying a map can equal one year of navigation experience for (global) determinations of Euclidean distance and relative object location. Conversely, knowledge gained through navigation is better for (local) determinations of orientation (angle from one object to another). Darken and Sibert [1996] performed compared user performance in a search and wayfinding task under conditions both with and without a map, showing the utility of the map in executing an efficient search strategy.

Orientation coupling is often performed during navigation through the real world: we often turn a map to match the direction we are facing at the moment. A display that mimics such map behavior is called a track-up or forward-up display. Levine et al. [1984] were the first to demonstrate experimentally the importance of map alignment to successful task completion. Later studies [Aretz 1991; Eley 1988] suggest that track-up displays are less confusing for novice users, but experts prefer the

north-up display as it matches the perspective of a remembered canonical view. North-up displays also have the advantage of providing a common frame of reference for communicating map information.

A further study by Darken and Cevik [1999] investigated the utility of orientation coupling in several search tasks. During the tasks, subjects were asked to navigate a virtual world with the aid of an overview map. The results confirm for these applied tasks what Aretz and Wickens [1992] had found in more abstract situations. They suggest that track-up map displays are best for tasks in which the search target is clearly marked on the map: the user can essentially navigate through just the track-up overview display. The results further suggest that north-up map displays are best for searches in which the destination is either unknown, or previously visited but not shown: the stability of the display of the overview map allows the user to remain oriented with respect to an organized search strategy.

The task we have developed for our study is the multi-perspective identification task. This task requires a subject to combine the information from two different perspectives to identify a target object. More specifically, a subject must select from among identical-looking objects on an overview map given distinguishing information in a local-perspective. This task might be considered the inverse of a virtual-world search in which the search target is marked on the map: rather than using an indication on the map display to guide oneself to the target, one must indicate on the map display where one sees the target from a local perspective. In other words, our task requires the user to use local view information to interpret a global map, whereas prior studies have required users to use map information to guide local actions.

In this paper, we explore the relative utility of view proxies, tethers, and track-up orientation coupling for helping a user to integrate multiscale information from multiple views. Our goal is to quantify the value of each linking method in terms of errors and decision time in performing a specific task. This task requires a global object-location decision to be made using an overview map, given local orientation in a vehicle view. We have run two experiments based on this task to evaluate the particular strengths each linking method contributes. Experiment 1 contrasts how users perform when directional proxies, tethers, and/or view couplings are available with a single local view. Experiment 2 addresses how much aid directional proxies and tethers give the user when two local views are present.

2. Shared Experimental Design

In order to quantify the value of each linking method, we developed the multi-perspective identification task. We ran two experiments based on one basic design. We describe the common aspects in this section, including the experimental task, the tether and proxy conditions, the apparatus, and the methods of measurement.

2.1 Multi-Perspective Identification Task

Consider a situation in which one or more autonomous vehicles are exploring a landscape, and it is the job of an operator at mission control to identify the position of certain objects spotted by one of the vehicles. As illustrated in Figure 2, the operator is given an overview map in one window and in sub-windows is given local forward-looking views from just above and behind each vehicle. The multi-perspective identification task consists of monitoring one or more local, vehicle-centric views for a

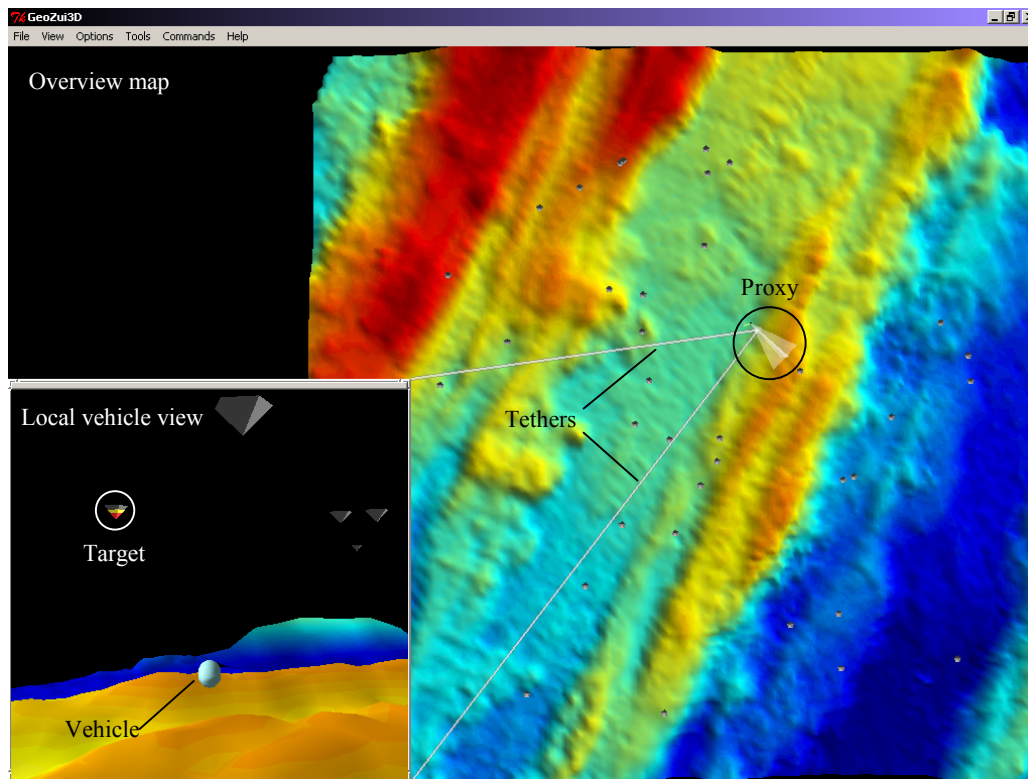


Figure 2: Example of the display presented to subjects in Experiment 1.

distinctive target within a field of distractors (Experiment 1 uses one local view, while Experiment 2 uses two). Each local view smoothly¹ follows a small vehicle as it wanders randomly through the environment. Once the target appears in a local view, the subject must identify it in the overview map by clicking on its representation with the mouse. An example of the display presented to users is shown in Figure 2. Subjects are told that they can make their decision based on any available information, including changes in vehicle heading, distractor layout, and surface cues such as form and color.

2.2 Initial Conditions Per Trial

Each trial starts with a new random layout of 35 distractors and a new random path for each vehicle to follow. The target is placed along the line of sight of a latter segment of one of the paths such that it would be encountered from a north-looking ($< 90^\circ$ from north) or south-looking ($< 90^\circ$ from south) vantage point. The path always starts near the center of the overview and continues in a constrained random manner. The path is a cubic spline created with 8 points (7 base segments). Each successive point is generated by selecting a random distance and a random direction. The distance is constrained to be between roughly 2-5% of the width of the screen from its predecessor, while the direction is constrained such that each base segment is within 90° degrees of its predecessor. If the path did not provide a way for the target to be encountered in the chosen direction for the trial, the path was regenerated from scratch.

¹ The vehicle maintains a heading tangential to a spline, as does the “camera” for the local view. The effect is that the vehicle always appears centered within the local view.

2.3 Linking Aids

The following conditions were common to both experiments:

- Proxy vs. no proxy
- Tethers vs. no tethers

The third linking method, using track-up orientation coupling, was only applied in the first experiment.

Regardless of the methods employed for a trial, a small (roughly 5-pixel-wide), semitransparent box appears on the overview map at the position of each vehicle. When tethers are present, they appear as semi-transparent lines connecting two corners of a local view to the center of the corresponding small box, as shown in Figure 1. When proxies are present, they appear as semi-transparent triangles (Figure 1). A small dot at the apex of the triangle represents the viewpoint for its corresponding local view.

2.4 Apparatus

Both experiments were run on a Windows 2000 (Professional) system configured with a Pentium 4 processor, a Wildcat II 5110 graphics card, and a 19” monitor running at a resolution of 1024x768 pixels. The animation rate of each experiment was roughly 30 frames per second. Subjects were provided a standard Microsoft mouse for controlling the on-screen cursor. All references in this paper to a click or a selection using the mouse refer to the clicking of the left mouse button only (input from the other two buttons were ignored by the experiment software).

2.5 Measurements

In both experiments, we measured decision time and errors. Decision time was measured as the amount of time that passed between when the target was first visible in a local view and when the subject moved the mouse cursor out of the local view to make

a selection on the overview map². Errors were recorded whenever the subject made an incorrect selection, and the position of that selection was recorded for later determination of the magnitude of the error (in angular degrees).

Other error conditions occurred, and trials ending in such conditions were recorded. It was possible for a subject to simply timeout by not responding to any target. If the target did not appear at all, the trial was repeated. Otherwise, the trial was recorded as an error (without a position), and the next trial was begun. Also, if the subject made an incorrect selection before the target ever appeared, the trial was repeated.

One last error condition was the result of an interface quirk we found necessary for the success of the experiment. In order for the subject to be able to make a selection in the orientation-coupled conditions, it was necessary to stop all vehicle movement when the subject was ready to make a selection. Movement was therefore stopped whenever the subject placed the cursor into the overview region, and subjects were instructed not to move the mouse cursor out of a local view until they were ready to make a selection. If the subject spent more than 5 seconds making a selection, or if the subject moved in and out of a local view too many times, the trial was recorded as an error (without a position), and the next trial was begun. At the end of each experiment, each subject was asked to fill out a short questionnaire. The questionnaire asked subjects to rate the usefulness of each linking aid encountered in the experiment. The scale went from 1 (counterproductive) to 5 (not useful) to 10 (extremely useful). The questionnaire also asked subjects which combination of aids they preferred the most.

3. Experiment 1

The first experiment compared user performance under all combinations of the three linking methods, with only a single local view present (as illustrated in Figure 2). The proxy and tether aids used in this experiment were described in the previous section; we now describe the orientation coupling used. When orientation coupling was enabled, the center of the overview tracked the vehicle's position; the center of the screen always corresponded to the position of the randomly wandering vehicle. In addition, the overview map was constantly rotated to track the forward direction of the vehicle; the "up" direction of the overview was always parallel to the vehicle's heading.

3.1 Subjects

Experiment 1 was run on 17 subjects: 10 male and 7 female.

3.2 Treatments and Design

The experiment presented each subject with a training block of 40 trials and 4 experimental blocks of 32 trials each. Each block was divided into sub-blocks of 8 trials each as shown in Figure 3. Subjects were asked to take a five to ten minute break between blocks 2 and 3.

The experiment was set up as a 2x2x2 within-subjects factorial design. Within each 2x2 sub-block, the presence or absence of *tethers* was varied, as was the expected *direction* that the local view would have to face for the target to be present (northern semicircle vs. southern semicircle). Each combination of these

variables appeared twice in each sub-block in a random order. Between sub-blocks, the presence or absence of the directional *proxy* was varied, as was the state of the orientation *coupling* (enabled vs. disabled). Sub-blocks were organized such that all conditions with the same value for *coupling* were grouped together, and that the order of each combination of *coupling* and *proxy* were counterbalanced across blocks.

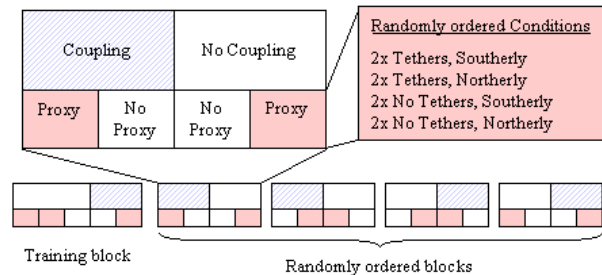


Figure 3: Experimental design of Experiment 1.

3.3 Results

We discarded all trials in which the subject failed to make a selection. This amounted to less than 3% of the total. The results for the remainder of the data are summarized in Figures 4 through 7.

3.3.1 Mean Error Rates

An analysis of variance on the mean error rate (within each subject) for each combination of *coupling*, *proxy*, and *tethers* revealed *coupling* and *proxy* to be highly significant. *Coupling* reduced errors by 27% ($F(1, 16) = 7.030$, $p < 0.001$), while use of a *proxy* reduced errors by 52% ($F(1, 16) = 33.097$, $p < 0.001$). The mean error rates for *proxy* and *coupling* conditions are summarized in Figure 4, in which the histogram bars indicate the estimated marginal means for each condition. The interaction between *proxy*, *coupling*, and *subject* was significant ($F(16, 16) = 3.423$, $p < 0.01$). This interaction effect is likely due to the sparsity of data in each cell (two samples) at this level of interaction. There was no main effect for *tethers* and no other interactions.

An analysis of variance was also performed on the mean error rate within each subject for each *direction*. This analysis indicated no main effect for *direction*.

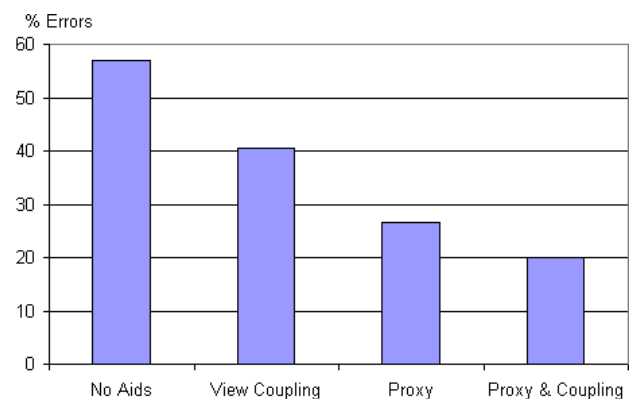


Figure 4: Percentage of errors made by subjects during Experiment 1 in each of the four linking categories tested.

² If the target was visible in a local view for less than a second, the timer was reset, under the assumption that the target moved too quickly across the view to be an appropriate starting point for measuring decision time. In practice, this rarely happened.

3.3.2 Error Magnitude

An analysis of variance was performed on *coupling*, *proxy*, *direction*, and *tethers* with respect to the magnitude of errors for all trials that ended in an error. Magnitude of error was measured in degrees as the angle between the ray from the vehicle to the target and the ray from the vehicle to the selected distractor (in the X-Y plane). The analysis again revealed *coupling* and *proxy* to be significant. *Coupling* reduced the magnitude of error by 38%, or 15° ($F(1, 45) = 7.468$, $p < 0.01$), while a *proxy* reduced error magnitude by 67%, or 26° ($F(1, 22) = 17.628$, $p \leq 0.001$). The interaction between *coupling* and *proxy* was significant as well ($F(1, 32) = 9.714$, $p < 0.01$), with the effects illustrated in Figure 5. The histogram bars in Figure 5 indicate the estimated marginal means among error trials for each condition of *coupling* and *proxy*. No other variables or combinations of variables were significant.

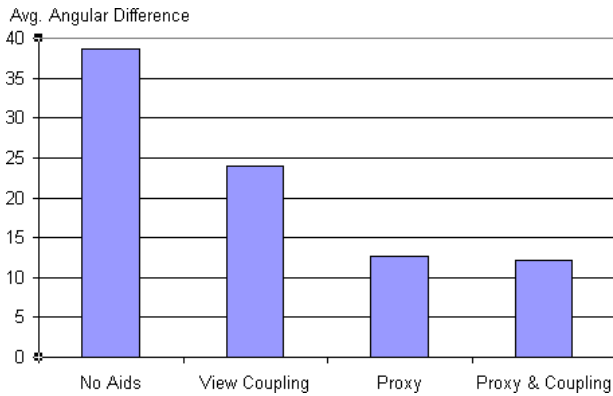


Figure 5: Average angular difference in degrees between the target and the subject's selection in Experiment 1.

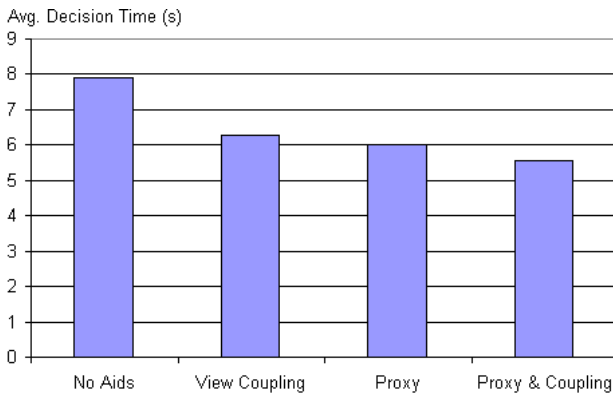


Figure 6: The average time elapsed in seconds from when the target first appeared in the local view and the subject's selection in Experiment 1.

3.3.3 Decision Time

We ran an analysis of variance on all four variables again with respect to decision time for all valid trials (ending in either success or error). As described in Section 2.3, decision time was measured as the time elapsed between the first appearance of the target and when the subject moved the cursor to make a selection on the overview map. This analysis showed that *coupling* reduced decision time by 15%, or 1 second ($F(1, 16) = 12.442$, $p < 0.01$), and a *proxy* reduced decision time by 18%, or 1.4 seconds ($F(1, 16) = 48.769$, $p < 0.001$). Their interaction was

also significant ($F(1, 32) = 13.505$, $p < 0.01$), as shown by the contrast of estimated marginal means in Figure 6. No other variables or combinations of variables were significant.

3.3.4 Subject Preference

Figure 7 summarizes the opinions of subjects collected in the questionnaire. The range of responses for both the *proxy* and *coupling* aids were between 7 and 10 (extremely useful), with averages of 9.1 and 8.4, respectively. The range of responses for the use of *tethers* was between 1 (counterproductive) and 7, with an average of 4.6. On the question of which combination of features was best, 8 answered with the *proxy/coupling* combination, 4 answered with all aids, and 1 answered with just the *proxy*. Four subjects misinterpreted the question, apparently answering instead which aid was most important. To this question, 2 answered with the *proxy*, and 2 answered with *coupling*.

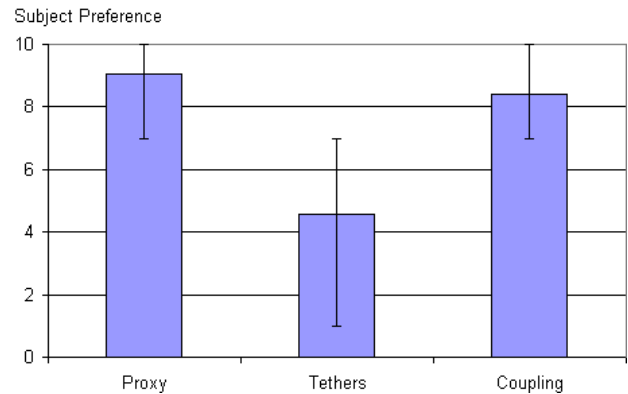


Figure 7: The average and range of subject opinions on the utility of each linking aid in Experiment 1.

3.4 Discussion

The results of Experiment 1 indicate that the *proxy* and *coupling* devices individually contribute to performance improvement in both accuracy and decision time, while the *tethers* do not. Our results further show that these two linking aids can be combined for further improvement in terms of reducing the number of errors. The results also indicate that the *direction* of the local view has no significant effect on accuracy or decision time. Subject preferences are in line with the performance data.

While *tethers* do not have a significant effect on this task with a single local view, we believed that *tethers* would benefit a task with multiple local views, at least when *coupling* was not in use. We believed this because adding another vehicle introduces ambiguity: Which view belongs to which *proxy*? This was our motivation for the second experiment.

4. Experiment 2

The second experiment compared user performance under combinations of only *tethers* and *proxies*, with two local views present (as illustrated in Figure 8). Each local view followed its own vehicle, and each vehicle followed its own constrained random path. Orientation *coupling* was inappropriate for this experiment because it would have greatly altered the nature of a subject's task execution. The purpose of this experiment was to determine whether or not *tethers* could be useful when ambiguity exists in the identity of the representations of independent views.

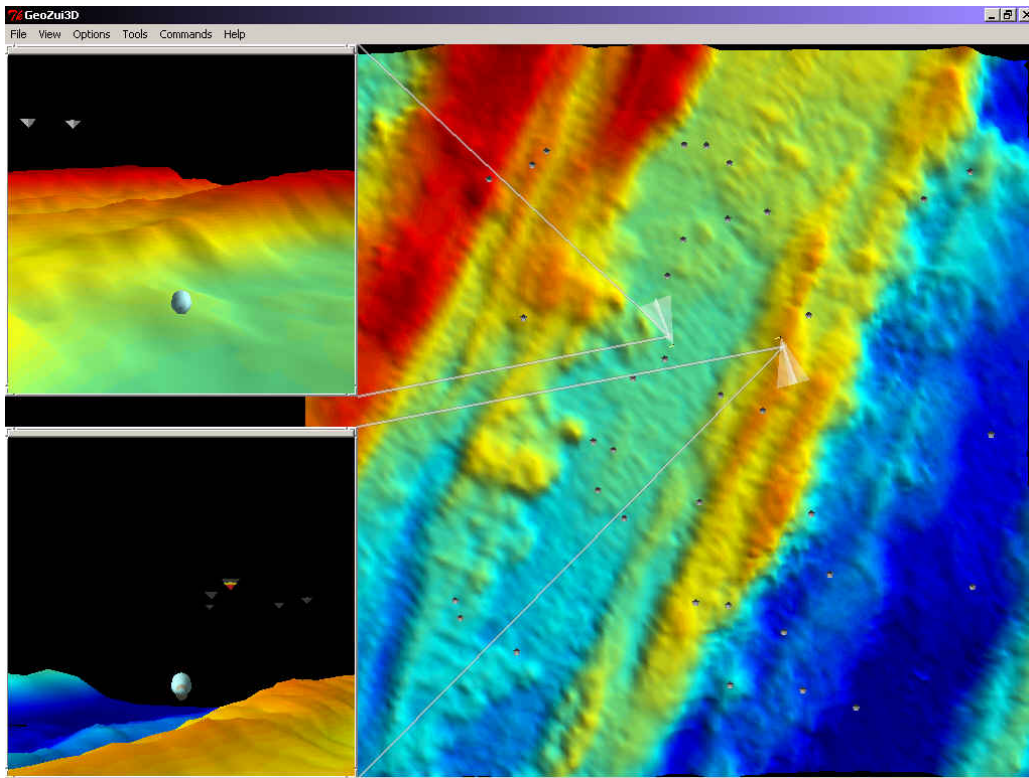


Figure 8: Example of the display in Experiment 2.

4.1 Subjects

Experiment 2 was run on 17 subjects: 11 male and 6 female.

4.2 Treatments and Design

The experiment presented each subject with a training block of 40 trials and 4 experimental blocks of 32 trials each. Each block was divided into sub-blocks of 8 trials each as shown in Figure 9. Subjects were asked to take a five to ten minute break between blocks 2 and 3.

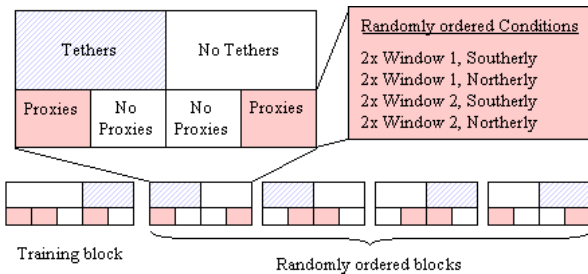


Figure 9: Organization of treatments in Experiment 2.

Within each sub-block, the *window* in which the target was expected to appear first was varied, as was the expected *direction* that the local view would have to face for the target to be present (northern semicircle vs. southern semicircle). Each combination of these variables appeared twice in each sub-block in a random order. Between sub-blocks, the presence or absence of the directional *proxies* was varied, as was the presence or absence of *tethers*. Sub-blocks were organized such that all conditions with the same value for *tethers* were grouped together, and that the order of each combination of *tethers* and *proxies* were counterbalanced across blocks.

4.3 Results

We discarded all trials in which the subject failed to make a selection. This amounted to about 5% of the total. We also discarded completely the results of one male subject; his comments and data indicated he did not even try to make a valid selection when there were no linking aids in place. The results for the remainder of the data are summarized in Figures 10 through 12 and Table 1.

4.3.1 Mean Error Rates

An analysis of variance on the mean error rate (within each subject) for each combination of *proxies*, and *tethers* revealed *proxies* to be significant. Use of *proxies* reduced errors by 47% ($F(1, 15) = 105.521$, $p < 0.001$), from an error rate of 69% to an error rate of 37%. The interaction between *proxies* and *subject* was significant ($F(15, 15) = 3.069$, $p < 0.05$). This interaction effect was in the amount that use of proxies helped; there were no subjects for which proxies degraded performance. There was no main effect for *tethers* and no other interactions.

A similar analysis of variance was performed on the mean error rate for each combination of *direction* and *window*. This analysis revealed that *direction* was significant: heading in a southerly direction when discovering the target decreased the error rate by 10%, from a rate of 56% to a 50% error rate ($F(1, 15) = 9.246$ and $p < 0.01$). There was no main effect for *window* and no interactions. Figure 10 illustrates the relative strength of *proxies* and *direction* using the estimated marginal means for each condition of their combination.

4.3.2 Error Magnitude

We performed an analysis of variance on *proxies*, *tethers*, *direction*, and *window* with respect to the angular difference for all trials that ended in an error. This analysis revealed only *proxies* to be significant. The presence of *proxies* reduced the average angle of error by 39%, from 51° to 31° ($F(1, 17) =$

14.938, $p < 0.01$). There were no main effects for any of the other variables, and no interactions.

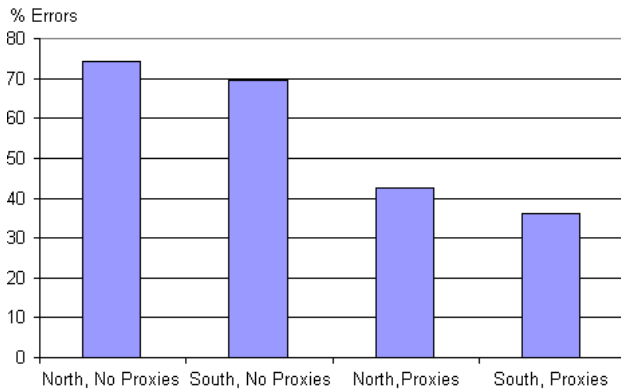


Figure 10: Percentage of errors made by subjects during Experiment 2 in each condition of *proxies* and *direction*.

4.3.3 Decision Time

An analysis of variance on the same four variables was performed with respect to decision time. This analysis found *proxies* and *window* to be the most significant individual factors. Presence of *proxies* decreased average decision time by 13%, from 13.9 seconds to 12.1 seconds ($F(1, 15) = 10.886$, $p < 0.01$). Having the target appear in the upper *window* as opposed to the lower decreased decision time by 11%, from 13.7 to 12.3 seconds ($F(1, 15) = 6.666$, $p \leq 0.05$). The relative contributions of *proxies* and *window* to decision time are shown in Figure 11, as indicated by the estimated marginal means for each condition.

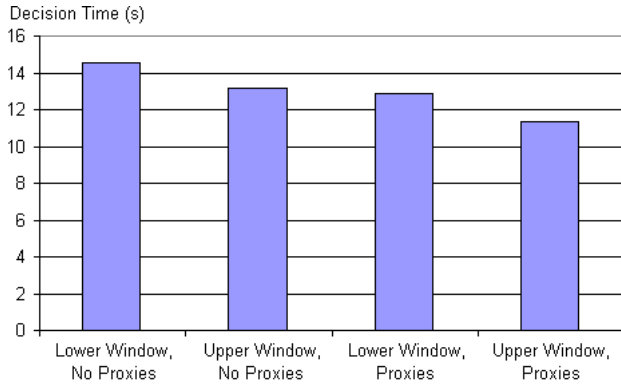


Figure 11: Average decision times in Experiment 2.

The presence of tethers reduced decision time by roughly 6%, from 13.4 seconds to 12.6 seconds. However, this failed to reach statistical significance ($F(1, 15) = 3.784$, $p \approx 0.07$). There was a significant interaction between *proxies*, *tethers*, and *direction* ($F(1, 16) = 5.667$, $p \leq 0.05$). The estimated marginal means for this interaction are given in Table 1. We have no explanation for this interaction.

4.3.4 Subject Preference

Figure 12 summarizes the opinions of subjects collected in the questionnaire. The range of responses for the aid of proxies was between 8 and 10 (extremely useful), with an average of 9.8. The range of responses for the use of tethers was between 1 (counterproductive) and 8, with an average of 6.6. On the question of which combination of features was best, 10 answered

with both proxies and tethers, and 6 answered with just the proxies.

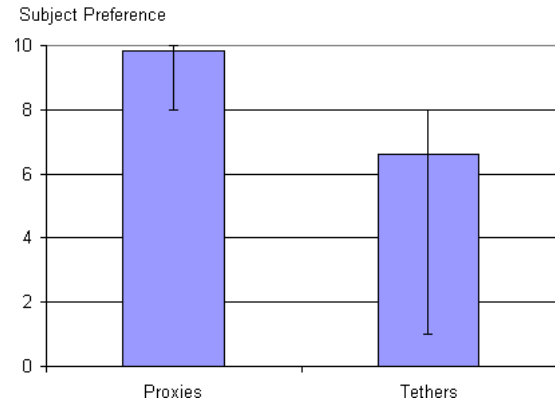


Figure 12: The average and range of subject opinions on the utility of each linking aid in Experiment 2.

	No <i>Proxies</i>		<i>Proxies</i>	
	No <i>Tethers</i>	<i>Tethers</i>	No <i>Tethers</i>	<i>Tethers</i>
Northerly <i>Direction</i>	14.1	13.4	11.2	11.5
Southerly <i>Direction</i>	14.4	13.6	13.9	11.8

Table 1: Interaction of variables on decision time. The highlighted cell indicates where the interaction is most apparent.

4.4 Discussion

The results of Experiment 2 indicate that the use of a *proxy* device contributes in both accuracy and decision time.

Although tethers failed to provide a statistically significant benefit, we feel that they might be useful especially in cases where there are many local views. We speculate that tethers might only be useful for a brief period of time, and that a better implementation might involve a strategy for having tethers appear briefly upon certain user actions.

In contradiction to the results of Experiment 1, *direction* had a small but significant effect on the number of errors. While we cannot be sure, this may have been due to the increased difficulty of Experiment 2 over Experiment 1, and a consequential increase of attention paid to the detail of the land underneath. Our observation of the terrain map indicates that the land to the south was more distinctive than the land to the north, and therefore may have provided better contextual cues.

Another result is that subjects responded faster when the target appeared in the upper *window* than in the lower window, by about 1.4 seconds. It appears that most subjects were paying attention primarily to the upper window, incurring a time penalty if the target appeared in the lower window. Four subjects appeared to pay attention primarily to the lower window. If the absolute values of differences between response times were taken across subjects, the average difference is much closer to 2 seconds, although the analysis does not indicate an interaction between *window* and *subjects*.

5. General Discussion

The strongest result from both experiments is the utility of directional proxies in reducing errors, with a reduction in error rates of around 50%. Orientation track-up coupling worked nearly as well for use with a single local view, but is not straightforward with multiple views. Tethers contributed little, even when two local views were present. These results suggest some guidelines for interface design, and provide inspiration for new possibilities.

1. *Use directional proxies.* Our results suggest that, for a multi-perspective task such as ours, a directional proxy should be made available for each local view. It appears that proxies such as ours aid in the mental transformation of an angle from a plane along the line of sight (the local view) to one perpendicular (the global overview). Our results do not indicate what precise characteristics a directional proxy should have, but we propose that the proxy should be as minimal as possible while still providing information about the extents of the viewing angle.

2. *Use track-up coupling or a similar aid to enhance the proxy.* Our results agree with and support Aretz [1991] and Eley [1988] in this matter. Track-up coupling simplified the experimental task, at least in part, by reducing the area of the overview that subjects had to consider—it was always between the middle and top of the screen in a relatively narrow area. Furthermore, if the target crossed directly in front of the local view, the subject could simply select the target in a straight line above the center of rotation (since there were rarely more targets along that line). Without a proxy, coupling does not provide guidance as to how to map an angle from the plane along the line of sight to one perpendicular. However, it does make consistent the notions of left and right within these two planes, and it does provide a line of reference for the forward direction. It would be interesting to compare track-up coupling with other aids. For instance, the same information might come across using a line that emanates from the eyepoint on the proxy in the forward direction, all the way across the screen. With such a device, the user could be reminded of which way left and right went by keeping a distinctive symbol to one side of the line.

Our results showed that tethers tended to add clutter without apparent benefit. Subjects were able to associate windows with proxies in their absence, presumably by using terrain matching and motion cues. However, this may have been partly due to the nature of our task, which gave subjects ample time to visually associate windows with proxies before the target appeared. It is quite possible that in a dual task situation, where an operator was using the three-dimensional display to monitor some situation intermittently, the tethers could be more useful. Also, it might be useful to display tethers only when the user needs them for making the decision as to which proxy (or moving object) belongs with which local view.

Alternatives to tethers include things such as color-coding the proxies to match the window borders of local views. Our results say nothing for or against the use of tethers in static images (such as magazine illustrations), nor in environments with many moving objects—these situations may still benefit from the constant use of tethers.

In the second experiment subjects tended to be faster when the target appeared in the upper of the two local view windows. This result suggests that when there are multiple local views, the interface should designate one as being primary, especially if there is a higher probability of needing information from that

view. This might be done by simple placement (for instance, the upper-left corner for use in western cultures), or by a distinctive border around the focal window.

6. Conclusion

Our experiments have quantified the utility of three devices for linking forward-looking perspective views with a plan-view window. The results show that both proxies and track-up coupling are effective devices for helping people to understand the relative spatial arrangement of the views. Counter to our expectations, tethers proved not to have a measurable benefit. However, this may have been partially due to the particular task constraints of our study.

The fact that tethers approached significance in the second experiment implies that there may be situations in which tethers could provide a benefit. One of the problems with tethers was that they introduced visual clutter, which suggests that some strategy for intermittent display might be beneficial. Possibilities for implementing just-in-time display of tethers include making the tethers appear only when the mouse cursor is over a proxy, or when the user actively selects the proxy (by clicking on it, for instance). It may also be useful to make the tethers appear for one or two seconds when the mouse cursor first enters a local view.

Track up coupling might also benefit from a more flexible approach. For instance, the track-up coupling could become active between a particular local view and the overview only when the user's cursor is within the local view. When the user exited, the coupling would be deactivated and the overview would either stop moving or animate back to a canonical orientation. We intend to investigate these ideas further.

7. Acknowledgements

This research was funded by NSF grant 0081292. We wish to thank Hannah Sussman for running the majority of subjects for these experiments.

8. References

- ARETZ, A. J. 1991. The Design of Electronic Map Displays. *Human Factors*, 33, 1, 85-101.
- ARETZ, A. J. AND WICKENS, C. D. 1992. The Mental Rotation of Map Displays. *Human Performance*, 5, 4, 303-328.
- DARKEN, R. P. AND CEVIK, H. 1999. Map Usage in Virtual Environments: Orientation Issues. *Proceedings of IEEE Virtual Reality 99*, 133-140.
- DARKEN, R. P. AND SIBERT, J. L. 1996. Wayfinding Strategies and Behaviors in Large Virtual Worlds. *Human Factors in Computing Systems CHI '96 Proceedings*, ACM Press/Addison Wesley, New York, 142-149.
- ELEY, M. G. Determining the Shape of Land Surfaces from Topographical Maps. 1988. *Ergonomics*, 31, 3, 355-376.
- LEVINE, M., MARCHON, I. AND HANLEY, G. 1984. The Placement and Misplacement of You-Are-Here Maps. *Environment and Behavior*, 16, 2, 139-157.
- MACKINLAY, J. D., ROBERTSON G. G., AND DELINE, R. 1994. Developing Calendar Visualizers for the

- Information Visualizer. *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST '94)*, ACM Press, New York, 109-118.
- RISCH, J. S., REX, D. B., DOWSON, S. T., WALTERS, T. B., MAY, R. A., AND MOON, B. D. 1997. The STARLIGHT Information Visualization System. *Proceedings of IEEE International Conference on Information Visualization*, London, 42-49.
- SEIGEL, A. W. AND WHITE, S. H. 1975. The Development of Spatial Representations of Large-Scale Environments. In H. W. Reese (ed.), *Advances in Child Development and Behaviour*, Academic Press, London, 9-55.
- STOAKLEY, R., CONWAY, M. J., AND PAUSCH, R. 1995. Virtual Reality on a WIM: Interactive Worlds in Miniature. *Human Factors in Computing Systems CHI '95 Proceedings*, ACM Press/Addison Wesley, New York, 265-272.
- THORNDYKE, P. W. AND HAYES-ROTH, B. 1982. Differences in Spatial Knowledge Acquired from Maps and Navigation. *Cognitive Psychology*, 14, 560-589.
- WARE, C. AND LEWIS, M. The DragMag Image Magnifier. 1995. *Human Factors in Computing Systems CHI '95 Companion*, ACM Press/Addison Wesley, New York, 407-408.
- YAMAASHI, K., COOPERSTOCK, J. R., NARINE, T., AND BUXTON, W. Beating the Limitations of Camera-Monitor Mediated Telepresence with Extra Eyes. 1996. *Human Factors in Computing Systems CHI '96 Proceedings*, ACM Press/Addison Wesley, New York, 50-57.

An Evaluation of Methods for Linking 3D Views: Plumlee, Ware

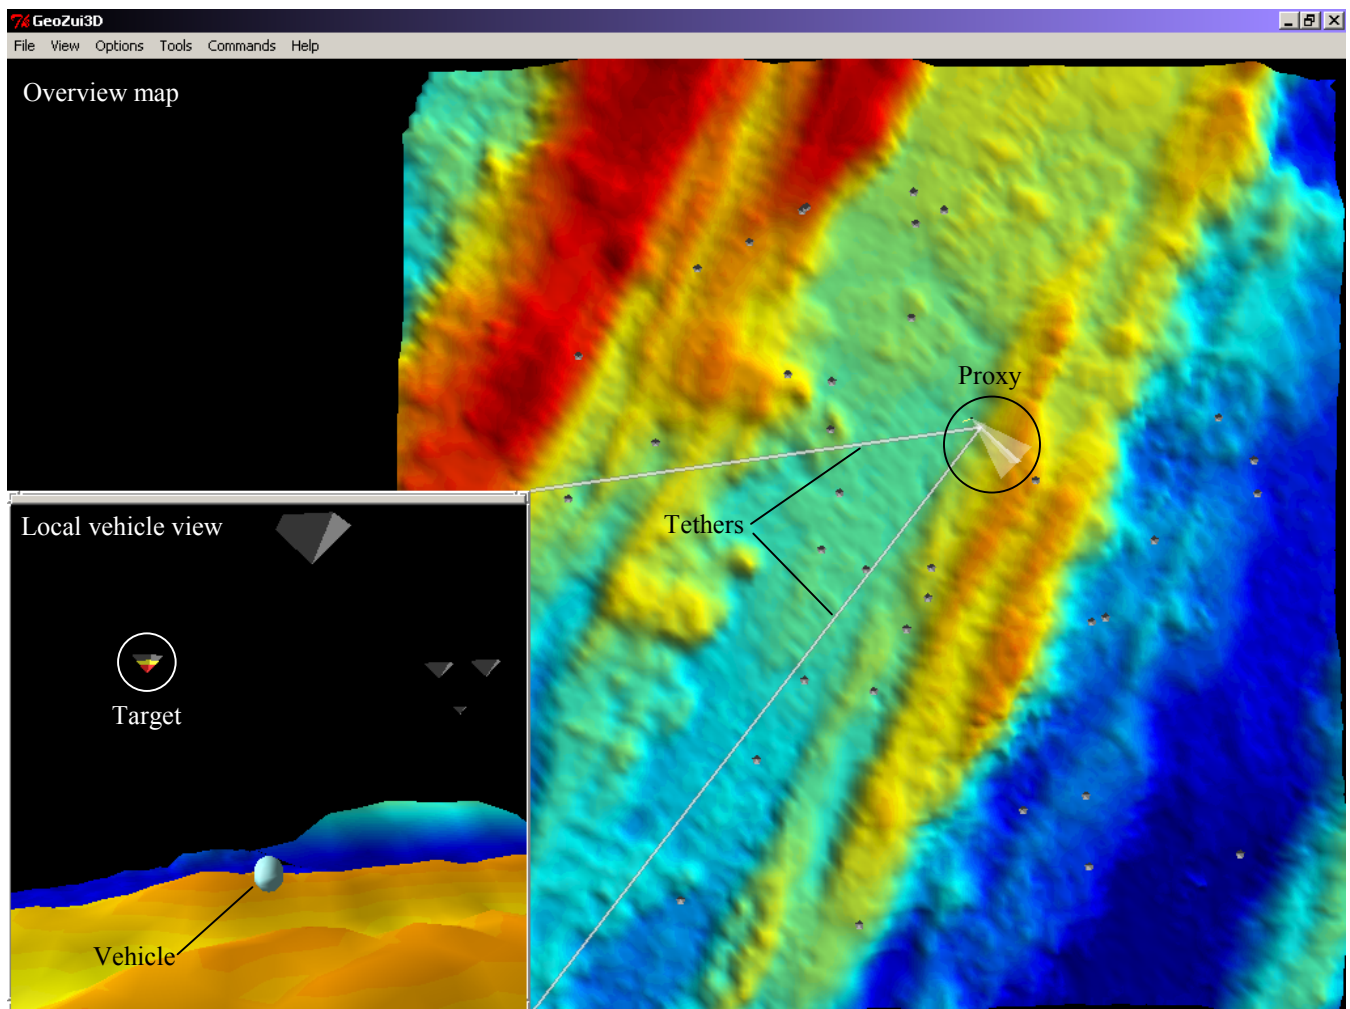


Figure 2: Example of the display presented to subjects in Experiment 1.