Pad vs. Leylines: Movement Model Matters

Susanne Jul
Electrical Engineering and Computer Science
University of Michigan
sjul@acm.org

ABSTRACT

This paper describes *Predictive Targeted Movement (PTM)* and reports on two experiments demonstrating its effectiveness. PTM is a constrained movement model that is defined in terms of cognitive concepts of navigation, such as location, destination and route. The exact definitions of these must be defined by the designer in accordance with the navigational needs of their design situation. PTM incorporates a notion of prediction that allows heuristic factors to be used in movement constraints. PTM has been applied to inter-object navigation in Jazz. Results from a study comparing PTM-based movement to a conventional movement model showed increased task performance, without increased error, and suggested that the physical and cognitive costs of navigation were reduced.

Keywords

Interaction Design, Navigation, Wayfinding, Locomotion, Predictive Targeted Movement, Multiscale, Zooming User Interface, Jazz, Lodestone, Leyline.

INTRODUCTION

Movement—changing location (at least conceptually)—is an important aspect of human-computer interaction. Users "go to" a web site, "enter" a chat room, "move" a file to the Trash, etc. Most movement in electronic environments is in service of *navigation*—determining where things are and getting to them—that is, it is a purposeful and directed activity [11]. Navigation, in turn, is in service of the user's task, e.g., gathering information about or editing the environment or its contents. Empirical evidence from the physical world indicates that the cognitive complexity of navigating in an environment is determined, in part, by what movement is possible in that environment [14, 15]. However, most efforts to support movement in electronic environments have focused on the physical cost [6, 7, 12] and on movement within a view [7, 12], disregarding cognitive considerations of movement design.

This paper describes *Predictive Targeted Movement (PTM)* and reports on two experiments demonstrating its effectiveness. PTM is a simple but powerful technique for designing movement in electronic spaces that combines the designer's knowledge of the user's task with the computational opportunities of electronic environments. PTM is a constrained movement model, that is, it limits, rather than increases, freedom of movement. PTM is defined in terms of cognitive concepts of navigation, such as location, destination and route, that depend on the specifics of a particular design situation.

PTM incorporates a notion of prediction that allows heuristic factors to be used in movement constraints. This element distinguishes it from other models that assume algebraic specifications of constraints [4, 5]. PTM differs from interaction techniques that constrain movement by computing trajectories automatically [7, 12], in that the user's destination need not be in the current view.

Prior efforts represent situation-specific movement models that encode knowledge of the user's task and the interaction environment. PTM, in contrast, is an abstract movement model. PTM relies on the designer to provide definitions of locations, routes and predictive functions that conform to the navigational needs of a particular design situation. Thus, the predictive element aside, PTM formalizes the intuitions underlying prior efforts. This level of abstraction permits PTM to be applied to a variety of design situations, and provides designers guidance in analyzing the navigational needs of a particular situation

This paper describes PTM, the application of PTM to inter-object navigation in Jazz [2], a spatial multiscale environment, and an experiment comparing PTM-based movement with a conventional design. The PTM-based movement model yielded a 30% reduction in time-on-task on a directed search task. This was accompanied by substantial and significant reductions in mouse activity both while directing movement and while, ostensibly, planning movement. Results suggest that the PTM-based design changes navigation in fundamental ways, and reduces both physical and cognitive efforts of navigation.

The PTM-based design and some of the experimental results were introduced briefly elsewhere [8, 9]. This paper details the PTM algorithm in full, illustrates its application to design, and elaborates and extends the empirical findings. Its contribution is to provide detailed empirical evidence for the effects of constrained movement on navigation as well as to elucidate PTM.

RELATED WORK

The present work distinguishes between movement and control of movement. While much work has focused on the latter, the present work focuses on how the navigational needs of the user's task can be used to determine what movement should be possible.

Other efforts have used knowledge of the user's task to define movement. Path Drawing [7] and Point of Interest Movement [12] both assume that the user's task is to move within a 3D virtual environment. Using Path Drawing, the system interprets a line drawn on a 2D view of space as a path in 3D space, using rules of physical movement, and moves the viewpoint along this trajectory. In Point of Interest Movement, the user indicates a point on an object and the system computes a trajectory to move the viewpoint to be centered on this point.

These approaches are similar to PTM in that they assume a target destination (a location or sequence of locations) and constrain movement to a path that leads to that target. However, they differ in that they assume knowledge of the nature of the target (a point or sequence of points in 3D space) and in that the user must specify the target explicitly. The latter limits their use to movement within a view.

Guided Navigation [4] and Constrained Navigation [5] also assume movement within a 3D space. Guided Navigation assumes that the user is touring the environment and constrains movement to follow a loosely scripted path, allowing the user some control of movement along and within this path. Constrained Navigation assumes that the user needs to examine individual objects or environmental features, and limits movement to regions defined parametrically with respect to these. While both constrain movement in accordance with the user's task, neither incorporates a notion of prediction and so do not allow constraints to incorporate heuristic considerations.

Of the work mentioned, only Path Drawing [7] and Constrained Navigation [5] report results of user studies. Of these, only Constrained Navigation examines the effect on any form of cognition, reporting on the effect on acquisition of spatial knowledge, but not on other aspects of navigational performance, such as measures of time and effort.

PREDICTIVE TARGETED MOVEMENT

Movement in PTM is relative to *lodestones* (so named because they exert navigational "pull")—locations the user is likely to want or need to go to in the course of performing their task. Movement is constrained to follow *leylines* (named for lines of power in Nordic and Celtic myths)—paths that lead from the current location to a lodestone. All movement must follow leylines and all leylines lead to lodestones.

Lodestones and leylines are defined by the designer in accordance with the needs of the user's task. Either may be defined explicitly or algorithmically. For instance, in a meteorological visualization tool, lodestones might be user-designated points of interest and/or locations of specific types of atmospheric conditions, e.g., the eye of a hurricane or leading fronts of masses of air. Leylines might follow pre-defined lines of interest, such as constant air-currents or contours of landmasses, or they might follow actual or predicted paths of weather systems. Definitions of lodestones and leylines can get as complex as appropriate for the task and might include nested definitions of lodestones or composite trajectories for leylines. Each might also incorporate dynamic elements such as the user's geographic or conceptual location.

Once the designer has decided what constitutes a lodestone and a leyline and how they are to be identified computationally, they must provide a predictive function that ranks a set of lodestones according to the likelihood of their being the user's desired destination at any given time. Prediction might be based on spatial proximity (e.g., to the mouse), relevance metrics in a query result set, past interaction, or other heuristic criteria. In the meteorological tool, lodestones representing weather phenomena might be ranked by severity or potential impact on human populations. The designer may also provide a predictive function for ranking leylines, if multiple leylines to a single lodestone are possible.

The system uses these functions to predict how likely it is that a given lodestone is the user's target destination and how likely it is that a given leyline is the user's desired route for getting there. Using these predictions, the system constrains movement to the most likely leyline. Optionally, the designer may provide a means for the user to refine predictions and negotiate the selected target and/or leyline. For example, a storm-watcher may be more interested in a weaker tropical storm about to make landfall than a more powerful hurricane still out to sea. The meteorological tool might provide a preview of high-probability destinations that includes distance from the current location, and allow the user to indicate a preference for the weaker storm.

The actual PTM algorithm is executed during interaction when the user signals their intent to move:

1. Identify the set of available lodestones
2. Rank the lodestones according to the predictive function, selecting the most likely destination(s)
3. [*Optional*] Present feedback to the user about the currently predicted destination(s) and negotiate the selection of alternatives
4. Compute or select a set of leylines to the predicted destination(s)
5. Rank the leylines according to the predictive function, selecting the most likely to be desired
6. [*Optional*] Present feedback to the user about the currently predicted leyline(s) and negotiate the selection of alternatives
7. Initiate and continue movement along the selected leyline until the target lodestone is reached

8. [*Optional*] If the inputs to the predictive function change during movement, repeat steps 2 - 7

In a simple example involving movement within a view, PTM might be applied to WIMP desktop interaction: Lodestones are icons and leylines are straight lines across the desktop. Prediction is based on mouse movement. The vector from the starting to the current position of the mouse indicates the general direction of the desired target. Lodestones within a thirty-degree arc are considered potential destinations and are ranked according to the deviation between the mouse vector and computed leylines. If there are no lodestones in the indicated direction, an area of the desktop is considered a virtual lodestone, allowing movement to uninhabited regions of the desktop.

Lodestones in the prediction set might be highlighted or magnified with the strength of the highlight or degree of magnification reflecting predicted probability. Alternatively, vectors indicating their respective leylines might be shown with length proportional to predicted probability. As the user moves the mouse, deviation from the most probable leyline causes the predictions to be updated. A “flick” of the mouse allows the user to approve the current target and complete the movement, selecting the target and moving the cursor, or sending a dragged item “skidding” onto the target. The user, when not dragging, might also approve the target simply by clicking at the present mouse location as though the target had been reached. The system would interpret this as a click on the actual target.

PTM IN JAZZ

PTM was applied to inter-object navigation in Jazz [2], a framework for designing and building multiscale electronic worlds using a zooming user interface. Like its predecessor, Pad++ [1], Jazz employs an interaction metaphor of a conceptually infinite two-dimensional surface that can be viewed at an infinite range of magnifications. Objects have position and extent on the surface. Their visibility can be configured to depend on the magnification (scale) of the view, e.g., becoming invisible when the amount of detail is too small to be useful. This metaphor implies a basic movement model of zooming (changing the view scale) that must be preserved if the purported benefits of multiscale environments [1] are to be preserved.

Lodestones and Leylines

A considerable amount of user interaction in Jazz entails inter-object navigation, that is, moving from one object to another. This task was adopted as the subject of an initial design study of the application of PTM. As the user’s goal is to move between objects, lodestones are, with one exception, defined to be individual objects, or rather, views at which an object is centered on the screen and appears at a “reasonable” magnification (here defined to be when it fills 90% of the view window along its largest dimension).

The inter-object navigation task implies that the user wants to move between objects as quickly as possible without needing exposure to intervening objects. Consequently, leylines should be the shortest paths possible. In multiscale, the shortest path between two objects is to zoom out until the objects appear to be one window-width apart, then zoom in on the target object [3]. However, arbitrary object layouts may not provide such trajectories naturally, i.e., there may not be an object to which to zoom out.

Figure 1 Lodestones and Leylines interaction in Jazz.

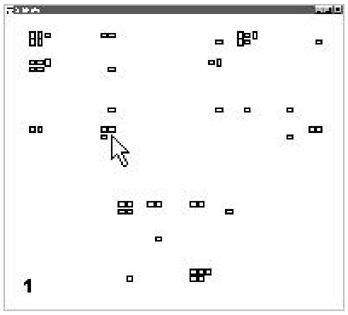

The user clicks on or near the object to which they want to go. The system predicts that the lodestone closest to the mouse is the intended destination.

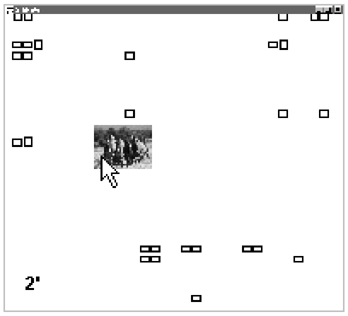

The system optionally shows feedback about the current prediction (here a thumbnail of the predicted target).

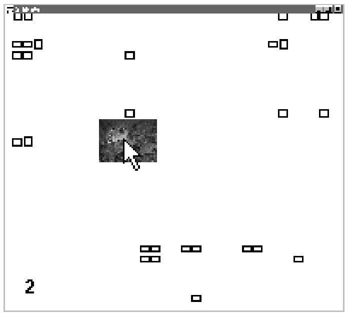

If the zoom-in destination prediction is incorrect, the user may correct it by moving the mouse without stopping the zoom. Once the prediction is correct, the user need not move the mouse again.

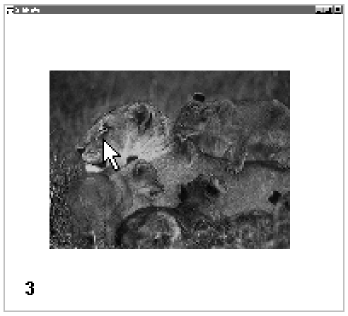

Zooming stops when the target is reached, if no further lodestones can be reached by zoom-in.

Thus, a special lodestone, the *Top of the World*, is introduced that guarantees the availability of zoom-out leylines. The Top of the World view is the most magnified view that contains all (other) lodestones in the world. The shortest path between two objects is approximated by zooming out toward the Top of the World, stopping when the target comes into the view, then zooming in on the target. Consequently, leylines are defined to be “straight” lines through space-scale. Movement between two objects, in many cases, entails following two leylines: a zoom-out followed by a zoom-in. Note that most leylines represent combined pan and zoom trajectories, that is, simultaneous movement in both planar and scale dimensions.

A more sophisticated design for this task would recognize the human predilection to use spatial proximity or visual grouping to convey semantic meaning, and include lodestones defined by spatial clustering of objects. This would introduce intermediate “Top of the Neighborhood” lodestones in both zoom directions. The simpler design, however, suffices to test the effect of movement model and minimizes confounding effects of specific design decisions.

Predictive Functions

In the absence of more specific knowledge of why the user is navigating between objects, target prediction is based on simple spatial proximity. The user indicates which direction in scale they would like to move, and the system predicts which lodestone is their intended destination. Zoom-out prediction is trivial, as the only lodestone considered is the Top of the World. Zoom-in prediction is based on mouse position. Disregarding the Top of the World, the system predicts that the lodestone closest to the mouse (in simple planar distance) is the desired destination. All leyline “prediction” is trivial, as there is only a single leyline from any given location to any given lodestone—there being only one “straight” line between two locations.

Because prediction is based on mouse position, target negotiation can be provided by monitoring mouse movement. When the system detects that the lodestone closest to the mouse differs from the current target, it updates the prediction and switches to the appropriate leyline.

Lodestones and Leylines Interaction

Figure 1 illustrates Lodestones and Leylines interaction. Because individual movements are constrained to follow leylines, overall movement is limited to the populated region of space-scale. This region is bounded by the Top of the World view and the space-scale extents of all objects, as shown in the space-scale diagram [3] in Figure 2.

Cognitively, Lodestones and Leylines interaction offers two potential benefits. The first benefit is to reduce the number of paths available to the user in any given view, as shown in the view schematics in Figure 2. Reducing the number of paths available in any given view reduces the complexity of individual decisions, and reduces the overall number of decisions users have to make during navigation. It also allows selection regions for each path in a view to be enlarged, simplifying movement control. The second benefit is the introduction of a fixed reference location, the Top of the World, to which users can return should they become lost or disoriented. These benefits are expected to reduce both the physical and cognitive costs of navigating.

Empirical Evaluation

The study comprised two experiments comparing the Lodestones and Leylines model of movement to a conventional model. Each experiment employed a 1 x 2 factorial within-subject design with repeated measures. The first factor, movement model, was manipulated within subject—all subjects used both models. The second factor, the order of presentation of the two models, was varied between subjects. The two experiments differed only in the amount of information the environment provided to aid navigation; the experimental tasks and movement support were identical.

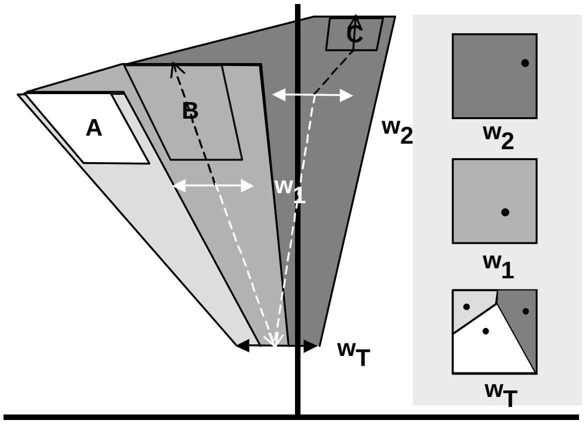

Figure 2 Space-scale diagram [3] of the movement constraints introduced by Lodestones and Leylines movement. Overall, movement is limited to the entire shaded region. Lodestones define sub-regions of movement limits. For instance, zoom-in anywhere in the dark gray region leads to lodestone **C** (e.g., by following the leyline indicated by the black dashed line from $\mathbf{w_2}$). Clicking to zoom out always leads to the Top of the World, $\mathbf{w_T}$, (e.g., by following the leylines indicated by the white dashed lines from $\mathbf{w_1}$, and $\mathbf{w_2}$). Schematics on the right show zoom-in sub-regions (movement options) at the indicated views, the dots indicating the positions of lodestones.

In the *Grid Markers* experiment, some navigational information was always available. This experiment was designed to simulate a task condition in which the user’s destination is not in the view, but where the user has an idea of where it is and the view presents enough information for them to get there. This is the normal case for much interaction, for example, getting to a file from the desktop by navigating through the file system, or getting to a particular piece of information in a website.

In the *Desert Fog* experiment, no navigational information was provided, except for object labels. This experiment was not intended to simulate a realistic task condition. “Desert fog” is a condition in which no information is available upon which navigational decisions can be based [10]. Although this situation is inherent to certain types of environments, including spatial multiscale, a realistic design would take steps to prevent it from occurring. This experiment tested the supposition that movement model can alter the demands of navigation so dramatically that an impossible task is made possible.

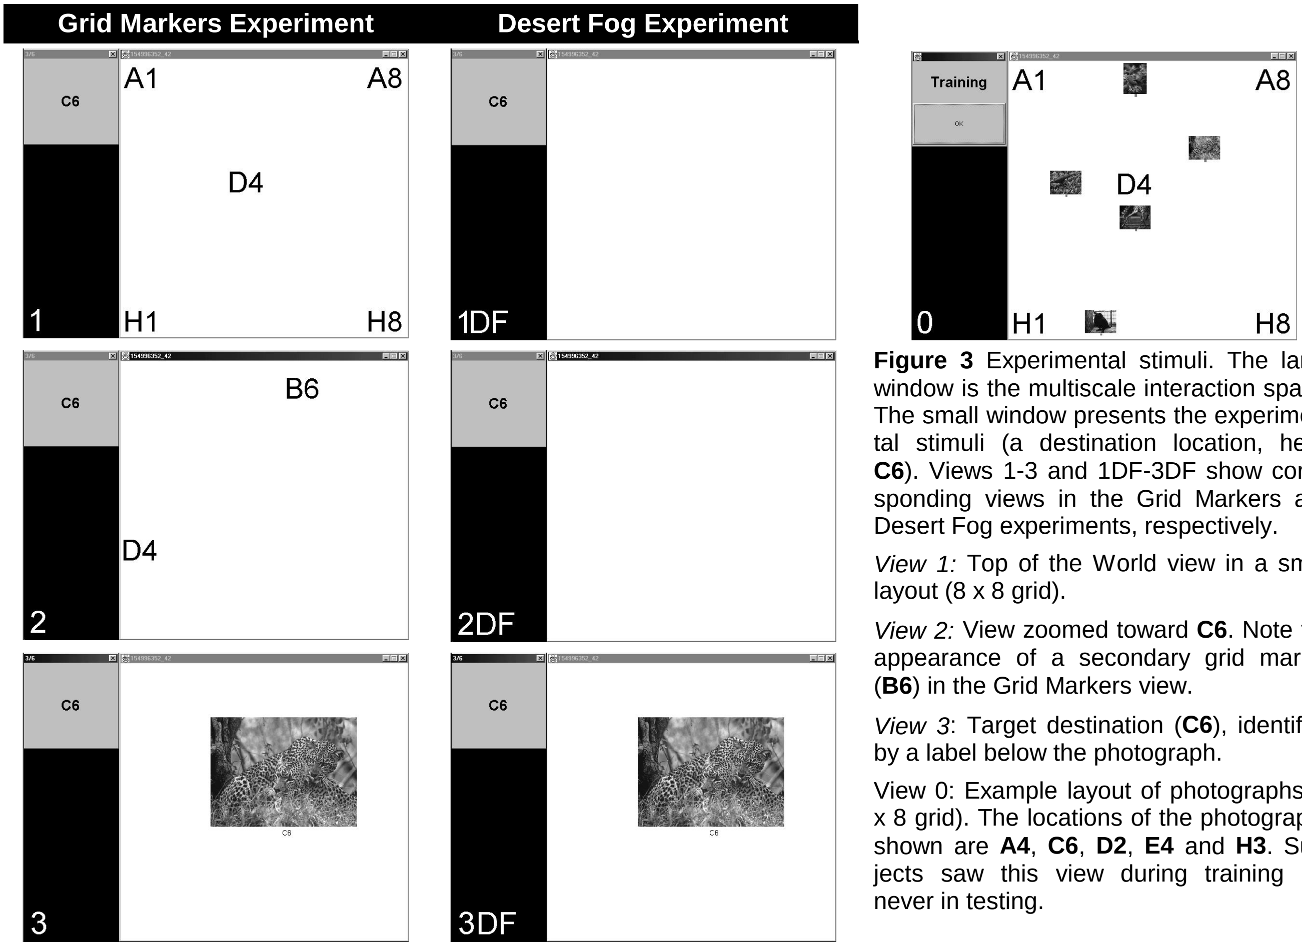

Figure 3 Experimental stimuli. The large window is the multiscale interaction space. The small window presents the experimental stimuli (a destination location, here, **C6**). Views 1-3 and 1DF-3DF show corresponding views in the Grid Markers and Desert Fog experiments, respectively.

View 1: Top of the World view in a small layout (8 x 8 grid).

View 2: View zoomed toward **C6**. Note the appearance of a secondary grid marker (**B6**) in the Grid Markers view.

View 3: Target destination (**C6**), identified by a label below the photograph.

View 0: Example layout of photographs (8 x 8 grid). The locations of the photographs shown are **A4**, **C6**, **D2**, **E4** and **H3**. Subjects saw this view during training but never in testing.

Subjects

25 subjects participated in the study, all volunteering in response to broadcast email. All were students or staff at the University of Michigan in disciplines ranging from music to computer science. 9 subjects were female, 15 were male, all between the ages of 18 and 50. All subjects reported at least one year of experience with mouse-based computers (only one less than three years) and average daily computer use of at least one hour a day. None reported prior familiarity with zooming user interfaces. Each subject participated in a single 1.5 – 2 hour session and was compensated with a $25 gift certificate.

Computational Environment

The study was conducted on a laptop computer with a Pentium II 266 MHz processor and 96 MB memory, running Windows 98, Java 1.3.1 and Jazz 1.0. The laptop's 12.1" display panel was used at a resolution of 800x600 pixels. The laptop's touch pad and mouse buttons were disabled and an ambidextrous external mouse used in their place.

Experimental Design

The two experiments were interleaved so that a subject performed both experiments with one movement model before repeating them with the other model. Performing the task in Desert Fog conditions requires full comprehension of the movement model in use, so the Grid Markers condition preceded the Desert Fog condition. Subjects were alternately assigned to start with one or the other movement model to counter-balance possible order effects.

Stimuli

The experimental stimuli are shown in Figure 3. Two fixed-size windows display the experimental cues and the Jazz interaction environment—small and large windows, respectively. The interaction environment consists of a set of photographs laid out on the surface. Photographs are selected randomly from a collection of professional photographs [13]; 50 in the Grid Markers experiment, 6 in Desert Fog.

Photographs all have the same size and aspect ratio, but may be in either portrait or landscape orientation. Their visibility parameters are configured so that a photograph is not visible until the view magnification is such that it covers at least 190 pixels along one dimension. This ensured that it was necessary to make navigational decisions with no photographs in the view, and that there was time to make such decisions while moving. Photographs all reach the visibility threshold at the same magnification and are spatially distributed to prevent visual occlusion (Figure 3, view 0).

The selection of photographs and their layout is random and unique to each training or testing run. Photographs are positioned within a conceptual grid. This grid is sized so that at most 10% of the cells will be occupied; a 23 x 23 grid in the Grid Markers experiment (50 photographs) and an 8 x 8 grid in Desert Fog (6 photographs). Photographs are placed in grid cells randomly with at most one per cell. The grid itself is not visible, but is the basis for a coordinate system for addressing locations. Rows are designated numerically, columns alphabetically (Figure 3, view 1).

In the Grid Markers experiment, the addresses of selected reference locations are displayed on the surface. Grid markers marking the four corners and the approximate center of the grid (Figure 3, view 1) are always visible (although they may not be contained in a given view, Figure 3, views 2-3). Secondary markers appear with each 1.75 increase in view magnification, ensuring that views contain at least one marker. Grid markers are fixed in size and do not change with view magnification. Grid markers are not available in the Desert Fog experiment (Figure 3, views 1DF-3DF), however, subjects were informed that each training or testing run started at the Top of the World. Each photograph is labeled with its grid address (Figure 3, views 3-3DF) in both experiments. No other location or target prediction feedback is provided.

Task

The experimental task is to move from one photograph to another. A random sequence of locations of photographs is selected without replacement, and presented, one at a time, to the user (Figure 3, small windows). Subjects move to the target location and press the space bar to indicate that they have arrived. If they are not at the correct location, their response is not accepted and they must continue to the correct location. If they are at the correct location, the next location cue is presented, and the subject seeks to go there from the present location. Subjects perform 15 and 5 trials (moving from one photograph to another) in the Grid Markers and Desert Fog experiments, respectively.

Movement Models and Control

The *Leylines* model of movement has already been described as Lodestones and Leylines interaction. The other movement model was the *Pad* model, which emulated the standard model offered by Pad++ [1]. In this model, movement is relative to the geometry of the space without regard for its contents. When zooming, whether in or out, the center of the zoom (around which the view expands or contracts) is the point on the surface on which the mouse was positioned when the zoom started. If the mouse is moved while zooming, the zoom center is the same surface point, but it and the entire surface move relative to the window. Panning, moving the surface relative to the window without zooming, also follows mouse movement.

A two-button mouse was used to control movement throughout the study. Both movement models used the left button to indicate zoom-out and the right to indicate zoom-in. In the Pad model, pressing the alt key (located symmetrically on either side of the space bar) and dragging with either mouse button pressed resulted in panning.

Procedure

All training and instructions were given by video and on-screen messages. An experimenter was present to answer subjects' questions during practice but not during testing sessions. Subjects received an introduction to the concepts of multiscale, zooming user interfaces, Jazz, and the experimental task before using either movement model.

After the introduction, they learned the first movement model to which they had been assigned, and practiced using it in the Grid Markers environment. First, they practiced with a small layout containing six photographs that were always visible (Figure 3, view 0). This allowed them to observe the behavior peculiar to the movement model. They then practiced on another small layout with normal visibility of photographs (Figure 3, view 1-3), and, finally, on a large layout, containing fifty photographs, like that used in testing in the Grid Markers experiment. Subjects were encouraged to practice as long as they liked, but were not allowed to stop until they had moved to five consecutive targets without error or appreciable hesitation.

After becoming familiar with the movement model, subjects were given tips on using it more effectively, e.g., moving the mouse during zoom-out to anticipate zoom-in. (Note that, by this point, most subjects had already discovered these techniques.) They then performed a final practice session on a large layout, and the Grid Markers test was administered. The Desert Fog experiment was then introduced and, after a single practice session, the Desert Fog test administered. Following a ten-minute break, the training and testing sequence was repeated using the second movement model to which they were assigned.

Data Collection

In addition to demographic and other individual information collected from each subject, behavioral data were recorded during their interaction with the software. This included the time spent on each trial—from the presentation of the location cue until the subject presses the space bar with that location in the view—and the number of response errors (i.e., spacebar pressed when the target is not in the view). View and mouse locations were sampled approximately every 100 ms. A sampling strategy was used rather than an event-driven record to avoid introducing different computational costs of data collection due to variations in event frequencies between the two movement models.

Computational Note

It should be noted that while the zoom rate—the change in magnification with each zoom increment—is constant and identical in both movement models, the computational overhead of zooming is somewhat larger in the Leylines model. In both models, the mouse location is sampled every 20 ms during zooming and, if it has changed, a system response computed. In the Leylines model, this causes the PTM algorithm to be executed. The PTM implementation was not optimized and an O(n) algorithm used for target

	Pad	Leylines	%	t(23)	p <
Time on task	94.2	66.0	-29.9	4.93	.0001
View move time (mouse press)	47.9	37.2	-22.3	3.12	.005
Mouse move time	52.4	27.9	-46.8	7.08	.0001
Mouse drag time	27.8	14.3	-54.4	5.22	.0001
Mouse non-drag time	24.7	13.6	-44.9	7.61	.0001

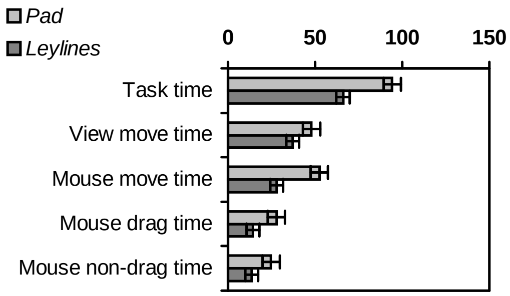

Table 1 Mean times per surface unit traveled. (milliseconds/surface unit)

	Pad	Leylines	%	t(23)	p <
View moves (clicks) (clicks/surface unit)	.055	.038	-30.9	4.72	.0001
Mouse moves (moves/surface unit)	.055	.049	-11.0	1.23	.25
Mouse drags (drags/surface unit)	.026	.029	11.5	1.21	.25
Mouse non-drags (non-drags/surface unit)	.029	.020	-31.0	3.47	.005

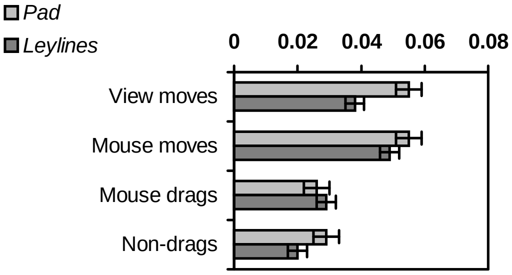

Table 2 Mean number of mouse actions per surface unit traveled.

Table 1 - Table 2 Time usage and mouse activity in Grid Markers condition. Measures are normalized to one net surface unit traveled. Regular type indicates values that are not statistically significant. % column shows change from Pad to Leylines.

prediction. (An O(log n) algorithm could be achieved by preprocessing the spatial layout of objects.) In the Pad model, a change in the mouse location causes a single translation of an affine transform.

Results

Data from one subject were eliminated due to faulty equipment discovered immediately following the session. Of the remaining subjects, 12 started with the Leylines model and 12 with the Pad model. Because the experiments each have only two conditions and Leylines is predicted to be superior, paired one-tailed t-tests were used.

Grid Markers Experiment

In order to understand the effects of movement model on task performance, physical effort and cognitive effort, the data are analyzed in terms of error, time on task(s), mouse activity, and relative use of time. Because of the randomness of the layouts and target sequences, certain measures of time and mouse activity are normalized to net planar distance traveled, that is, the total planar distance between targets in a given target sequence. This distance is proportional to the length of the shortest paths through space-scale. Planar distance is measured in surface units, which, at the canonical magnification of 1, correspond to pixels.

Error

There was no significant difference in the number of times subjects pressed the space bar erroneously, t(23) = .24, p < .6.

Time

The results for time spent on various tasks are shown in Table 1. All but three subjects were faster overall when using the Leylines model, moving the same distance in 30% less time, t(23) = 4.93, p < .0001. They also spent less time moving the view (*view move time*), 22% less in Leylines than in Pad, t(23) = 3.12, p < .005.

Subjects also spent less time moving the mouse (*mouse move time*) in Leylines, 47% less than in Pad, t(23) = 7.08, p < .0001. Examined more closely, mouse move time is divided into *drag* and *non-drag* time, time spent moving the mouse with and without a button pressed, respectively. Both were smaller in the Leylines condition. Drag time was 49% less, t(23) = 5.22, p < .0001, and non-drag time 45% less, t(23) = 7.61, p < .0001. (Note that mouse drag time is a subset of view move time as the view always moves when a mouse button is pressed, regardless of whether the mouse is moving.)

In short, overall time on task and analyzed subtasks was significantly reduced by the Leylines technology.

Mouse Activity

Mouse activity is an indicator of the physical effort expended. It is measured in terms of number of mouse actions (button presses and mouse moves), duration (in time) of actions, and distance the mouse is moved. A *view move* is synonymous with a mouse button press. A *mouse move* is a sequence of mouse position samplings in which the position changes at least every 150 ms. (This threshold is necessary to eliminate false “stops” introduced by computational de-

	Pad	Leylines	%	t(23)	p <
View move (mouse press)	877.6	1007.1	14.8	2.83	.01
Mouse drag	1126.6	474.5	-57.9	6.00	.0001
Mouse non-drag	932.5	679.2	27.2	5.62	.0001

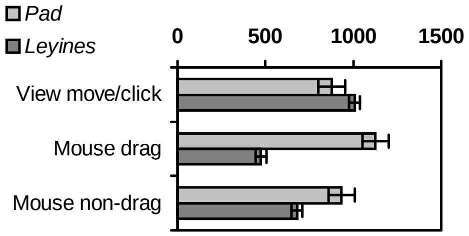

Table 3 Mean durations of mouse actions. (milliseconds)

	Pad	Leylines	%	t(23)	p <
Mouse drag	474.5	75.8	-84.0	5.26	.0001
Mouse non-drag	276.3	135.5	-51.0	5.31	.0001

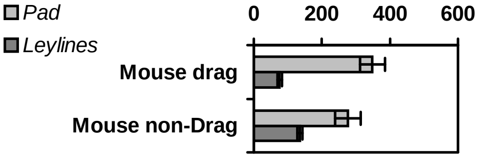

Table 4 Mean distances of mouse movement. (pixels)

Table 3 - Table 4 Mean durations and distances of mouse actions in Grid Markers experiment. % column shows change from Pad to Leylines.

lays, e.g., for garbage collection, and was determined by experimentation with mouse sampling.)

The results for the number of mouse actions are shown in Table 2. Subjects moved the view (pressed a mouse button) 31% fewer times per unit traveled in the Leylines condition, $t(23) = 4.72$, $p < .0001$. They also moved the mouse 10% fewer times, but this was not statistically significant, $t(23) = 1.23$, $p < .25$. Examined more closely, mouse moves are divided into *drags* and *non-drags*, mouse movement with and without a button pressed, respectively. (Mouse drags are, of course, a subset of view moves.) Subjects dragged the mouse 12% more times in Leylines, this was not statistically significant, $t(23) = 1.21$, $p < .25$. They non-dragged the mouse 30% fewer times in the Leylines condition, $t(23) = 3.47$, $p < .005$.

The results for durations and distances of mouse actions are shown in Table 3 and Table 4, respectively. Each view move (mouse press), on average, lasted 15% longer in the Leylines condition, $t(23) = 2.83$, $p < .01$. The average mouse drag was 58% shorter in time in Leylines, $t(23) = 6$, $p < .0001$, and 78% shorter in distance, $t(23) = 5.26$, $p < .0001$. The average mouse non-drag was 27% shorter in time in Leylines, $t(23) = 5.62$, $p < .0001$, and 51% shorter in distance, $t(23) = 5.31$, $p < .0001$.

In short, when using the Leylines model, subjects moved the view (pressed a mouse button) fewer times, but each move was longer in time. The number of times they moved the mouse was not significantly different. The number of times they moved the mouse with a button pressed (i.e., the view was moving) was also not significantly different, but each move was substantially shorter in both duration and distance. When a mouse button was not pressed (i.e., the view was stationary), subjects moved the mouse less often and moves were shorter in both duration and distance.

Relative Use of Time on Subtasks

In order to examine whether and how cognition is affected, the distribution of time on subtasks is examined. This analysis reveals whether the movement technology affected how subjects used their time. The relationships examined are distribution of view and mouse movement within the overall task, distribution of mouse movement subtasks within overall mouse movement, and distribution of mouse movement subtasks within view movement subtasks. These results are shown in Table 5.

With the Leylines model, subjects spent 42% of the overall task time looking at a stationary view, whereas they spent

	Pad	Leylines	%	t(23)	p <
View non-move time/ Time on task	.49	.42	-14.3	3.89	.001
Mouse move time /Time on task	.57	.43	-24.6	6.84	.0001
Mouse drag time/ Mouse move time	.52	.50	-3.8	.53	.6
Mouse drag time/ View move time	.58	.38	-34.5	7.24	.0001
Mouse non-drag time/ View non-move time	.56	.48	-9.0	3.62	.005

Table 5 Proportion of time on task spent on subtasks. % column shows change from Pad to Leylines.

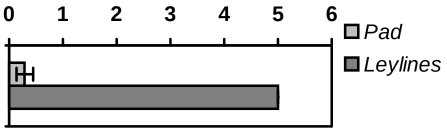

Figure 4 Mean number of trials (out of 5) completed in the Desert Fog condition. t(23) = 31.57, p = 0.

49% of their time doing so with the Pad model, t(23) = 3.89, p < .001. They spent 43% of the overall time moving the mouse with Leylines, but 57% doing so with Pad, t(23) =6.84, p < .0001.

Distribution of mouse movement subtasks was approximately equal in both models, with subjects spending 50% and 52% of the total mouse move time in dragging with Leylines and Pad, respectively, t(23) = .53, p < .6.

While moving the view, subjects spent 38% of the time also moving the mouse (dragging) with Leylines, and 58% doing so with Pad, t(23) = 7.24, p < .0001. While the view was stationary, subjects spent 48% of their time moving the mouse (non-dragging) in the Leylines condition and 56% doing so in the Pad condition, t(23) =3.62, p < .005.

In short, with the Leylines model, subjects spent a smaller percentage of their time looking at a stationary view and a smaller percentage of their time moving the mouse. They also spent a smaller percentage of time moving the mouse, both while moving the view and while looking at a stationary view. There was no significant difference in the percentage of time spent dragging and non-dragging when moving the mouse (disregarding whether the view was also moving).

Desert Fog Experiment

The only data analyzed in the Desert Fog experiment were trial completions. With the Leylines model, all subjects completed all five trials successfully. No subject was able to complete all trials using the Pad model, giving up (discontinuing the run) after an average of .29 trials, t(23) = 31.57, p = 0 (Figure 4).

Qualitative Results

In a post-test questionnaire, subjects reported greater satisfaction with the Leylines movement model. 16 of the 24 subjects stated that they preferred or strongly preferred Leylines, in general. When asked which model they would prefer if they “were doing something else while [they] were performing this task—say talking on the phone,” 21 subjects favored Leylines. 19 subjects found Leylines easier or much easier to use, while 3 thought Pad was easier, and 2 subjects thought they were about the same.

Many subjects cited the ability to return to the Top of the World and the need for less accurate mouse control as particularly attractive features of Leylines. Many cited the ability to pan as a positive feature of the Pad model and lack thereof a defect of Leylines.

Discussion

Results of the Grid Markers experiment show that the PTM-based technology, Leylines, increased task performance as measured by time, without increasing error. At the same time, the physical effort required to perform the task was reduced, as shown by the reductions in number and size of mouse actions. These results are straightforward.

Other results are more subtle. Subjects spent less total time, in the PTM-based design, looking at stationary views, implying that less time was dedicated solely to planning movement. That they also spent a smaller proportion of time looking at stationary views indicates that they were able plan movement faster or were able to do more planning while moving. The latter could result from the transfer of cognitive resources from movement control to movement planning, permitted by the reduced physical cost of moving.

The decreased mouse activity when not moving indicates that subjects were less confused or less agitated with the PTM-based technology. Anecdotal evidence including explicit comments about being lost or confused, and observed patterns of non-drag mouse movement (tracing out grid references and agitated “doodling”) supports the supposition that users were more confused with the conventional model and that this was, at least in part, due to spatial disorientation. That users felt better spatially oriented during movement with the PTM-based model is affirmed by the fewer but longer view moves. These suggest that users had more confidence in their plans and required less “stop and go” movement to review or adjust them. Interestingly, subjects used less time to move the same distance (recall that zoom speed was constant), indicating that they were following closer-to-optimal paths, despite devoting less time solely to movement planning.

Results from the Desert Fog experiment show that the PTM-based technology changes the navigational task fundamentally. All users eventually lost both spatial orientation and knowledge of productive actions in the conventional design. The PTM-based technology permits the user to become spatially disorientated, but provides a default action for reorientation, namely zooming out to the Top of the World. However, many users were unable to locate even the first target with the conventional design, although they, knowingly, started at the Top of the World. This suggests that the PTM-based technology reduced the need for maintaining spatial orientation, and the decrease in experienced disorientation in the Grid Markers experiment may have been due to decreased need rather than increased certainty or better maintenance of orientation during movement.

While the reduced freedom of movement would account for some of these effects, the predictive element of the PTM-based design must also play a role. Without the heuristic of spatial proximity, the user would have to select the desired target precisely in order to select its leyline, making Desert Fog navigation considerably more difficult. However, the present data are insufficient to dissociate the contributions of these two factors.

FUTURE WORK

Future work includes experimentation with more sophisticated definitions of lodestones, leylines and predictive functions. Lodestone experimentation is planned to consider navigationally significant locations, such as locations containing grid markers, and spatial clustering of objects. Leylines experimentation is considered to maximize the number of lodestones in view at any given time and follow common spatial structures. Prediction enhancement includes consideration of dynamic factors, e.g., disregarding the most recent lodestone visited, and adapting to patterns of movement by favoring frequently visited lodestones.

A different direction of future work is to explore ways of increasing the effectiveness of the Lodestones and Leylines design by experimenting with different types of prediction feedback and negotiation. Most importantly, PTM is expected to be applied and tested in other design situations, including implementing the desktop design outlined here and other designs that have been developed.

CONCLUSIONS

Predictive Targeted Movement is an abstract movement model that compels designers to consider movement in terms of elements of the user's task and the interaction environment. Movement is defined relative to lodestones—potential destinations in the user's task or navigationally significant locations—and is constrained to follow leylines—paths to lodestones that conform to the user's task or navigational needs. Using a predictive function supplied by the designer, the system eliminates unlikely navigational options, simplifying navigational decision-making. Applying PTM to inter-object navigation in Jazz yielded a design that, in empirical testing using a directed search task, increased task performance, without increased error, while reducing physical and cognitive effort.

The potential benefit of PTM lies in suggesting and guiding consideration of high-level cognitive issues during design of low-level interaction. Whether this benefit can be realized will depend on the difficulty of determining useful definitions of lodestones and leylines, and the feasibility of more sophisticated predictive functions. The immediate benefit, however, lies in demonstrating that profound cognitive effects can be achieved, without sophisticated technological solutions, by careful consideration of cognitive implications of low-level interaction design.

ACKNOWLEDGEMENTS

Special thanks to George Furnas for intellectual support, Mike Zyda and Dan Russell for video production facilities, and participants of the studies for their time and patience.

REFERENCES

1. Bedersen, B. B., Hollan, J. D. (1994). Pad++: A Zooming Graphical Interface for Exploring Alternate Interface Physics. *Proceedings of ACM UIST'94*, ACM Press, 17-26.
2. Bederson, B., Meyer, J., Good, L. (2000). Jazz: An Extensible Zoomable User Interface Graphics Toolkit in Java. *Proceedings of ACM UIST 2000*, ACM Press.
3. Furnas, G. W., Bederson, B. B. (1995). Space-Scale Diagrams: Understanding Multiscale Interfaces. *Human Factors in Computing Systems CHI ' 95 Conference Proceedings*, vol. 1, ACM Press, 234-241.
4. Galyean, T. A. (1995). Guided Navigation of Virtual Environments. *1995 Symposium on Interactive 3D Graphics.* ACM Press. 103-104.
5. Hanson, A. J., Wernert, E. A., Hughes, S. B. (1999). Constrained Navigation Environments. *Scientific Visualization: Dagstuhl ' 97 Proceedings*, IEEE Computer Society Press. 95-104.
6. Hinckley, K., Pausch, R., Goble, J. C., Kassel, N. F. (1994). A Survey of Design Issues in Spatial Input. *Seventh Annual Symposium on User Interface Software and Technology.* ACM Press. 213-222.
7. Igarashi, T., Kadobayashi, R., Mase, K., Tanaka, H. (1998). Path Drawing for 3D Walkthrough. *Proceedings of ACM UIST 98*, ACM Press, 173-174.
8. Jul, S. (In Press). "This is a *Lot* Easier!": Constrained Movement Speeds Navigation. *ACM Conference on Human-Factors in Computing Systems, CHI 2003.* Short paper.
9. Jul, S. (2002). Predictive Targeted Movement in Electronic Spaces. *ACM Conference on Human-Factors in Computing Systems, CHI 2002.* Poster.
10. Jul, S., Furnas, G. W. (1998). Critical Zones in Desert Fog: Aids to Multiscale Navigation. *ACM Symposium on User Interface Software and Technology, UIST 98.*
11. Jul, S., Furnas, G. W. (1997). Navigation in Electronic Worlds. *SIGCHI Bulletin*, 29, 4, 44-49.
12. Mackinlay, J. D., Card, S. K., Robertson, G. G. (1990). Rapid Controlled Movement Through a Virtual 3D Workspace. *SIGGRAPH '90 Conference Proceedings*, in *Computer Graphics* 24 (4), 171-176.
13. National Geographic Society. (1997). *National Geographic Photo Gallery.* CD-ROM.
14. Passini, R. (1996). Wayfinding Design: Logic, Application and Some Thoughts on Universality. *Design Studies*, Vol. 17. Elsevier Science. 319-331.
15. Peponis, J., Zimring, C. M., Choi, Y. K. (1990). Finding the Building in Wayfinding. *Environment and Behavior*, 22 (5). 555-590.