Can Smooth View Transitions Facilitate Perceptual Constancy in Node-Link Diagrams?

Maruthappan Shanmugasundaram*
Department of Computer Science
University of Manitoba

Pourang Irani†
Department of Computer Science
University of Manitoba

Carl Gutwin‡
Department of Computer Science
University of Saskatchewan

ABSTRACT

Many visualizations use smoothly animated transitions to help the user interact with information structures. These transitions are intended to preserve perceptual constancy during viewpoint transformations. However, animated transitions also have costs – they increase the transition time, and they can be complicated to implement – and it is not clear whether the benefits of smooth transitions outweigh the costs. In order to quantify these benefits, we carried out two experiments that explore the effects of smooth transitions. In the first study, subjects were asked to determine whether graph nodes were connected, and navigated the graph either with or without smooth scene transitions. In the second study, participants were asked to identify the overall structure of a tree after navigating the tree through a viewport that either did or did not use smooth transitions for view changes. The results of both experiments show that smooth transitions can have dramatic benefits for user performance – for example, participants in smooth transition conditions made half the errors of the discrete-movement conditions. In addition, short transitions were found to be as effective as long ones, suggesting that some of the costs of animations can be avoided. These studies give empirical evidence on the benefits of smooth transitions, and provide guidelines about when designers should use them in visualization systems.

Keywords: smooth transitions, animations, node-link connectivity, structure recognition, evaluation.

1 INTRODUCTION

Many recent visualization systems implement smoothly-animated transitions when shifting between different views of a visual structure, including transformations such as navigation, rotation, hiding and revealing structure, zooming in and out of the space, or switching between detail view and overview. In the visualization literature, the term *perceptual constancy* was introduced in [11] to suggest that smooth transitions can help the user maintain a sense of the true nature of the information despite the visual changes that occur during view transformations. Designers believe that smooth transitions will result in reduced time and effort as users mentally reorient themselves to the structures visible at the completion of the transformation.

Although smooth transitions have become a component in many visualizations there is little empirical evidence about whether smooth transitions really do facilitate perceptual constancy in viewpoint changes. While intuition suggests that smooth transitions may reduce cognitive load, there is also evidence that the time delays caused by animations can be disruptive, reduce efficiency and lead to frustrations [14]. Therefore, it is important to understand whether the use of transitions in visualization systems is effective.

We define a *transition* as a shift in the visual display from one view to another – that is, at time t_s (start time) the visualization presents view v_s (view at start time) and at time t_f (final time, $t_f > t_s$) the visualization presents view v_f (final view). A *smooth transition* is one that presents a number of intermediate frames or views (v_i) between t_s and t_f (Figure 1). Typically a minimum amount of geometric interpolation is necessary to shift between views v_s and v_f. This definition implies that smooth transitions have a direction of movement and occur at a defined speed.

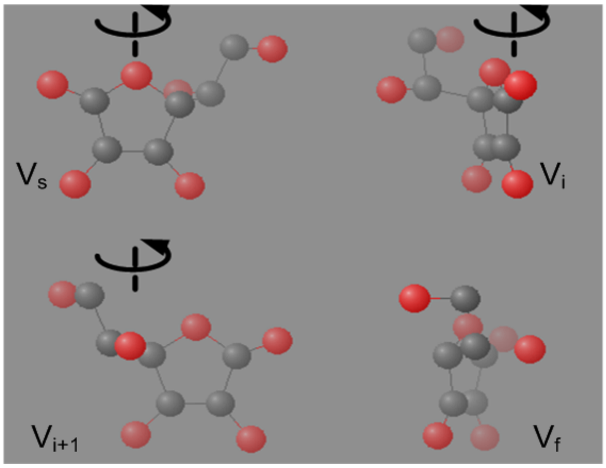

Figure 1 – Node-link structure of a chemical compound under smooth (counter-clockwise) rotation from v_s to v_f, allowing chemists to understand the structure of the molecular bonds.

The domain we use to study the effects of smooth transitions is node-link diagrams – in particular, graphs where the entire structure cannot be seen all at once in the viewport (either due to occlusion, or to the size of the graph). In these systems, there are two tasks that are strongly related to perceptual constancy: perception of connectivity, and perception of overall structure. First, the user's perception of connected elements can be disrupted as they move their view, and if the visual presentation impairs the perception of connections, then perceptual constancy is weakened. Second, users should be able to build and maintain a correct knowledge of the overall structure of the graph during viewpoint changes. If this condition is not fulfilled, users have to internally reorganize the structural elements of the node-link diagram, adversely affecting perceptual constancy.

The experiments described below use these two tasks to investigate the effectiveness of smooth transitions. If smooth transitions help to provide perceptual constancy as users move around in node-link diagrams, then users should perform better in tasks that require understanding of connectivity and overall structure. In addition, the studies also look at the issue of animation delay – that is, whether longer transitions better maintain perceptual constancy than short transitions.

The studies showed that smooth view transitions do have a beneficial effect on users' understanding of connectivity and structure, and that the effects are substantial. Errors were reduced by almost half when transitions were used, and these participants also moved their viewports significantly less often, and found the correct answers in significantly less time. Furthermore, we found that short transitions are often just as effective as long ones,

---

*e-mail: sundaram@cs.umanitoba.ca
†e-mail: irani@cs.umanitoba.ca
‡e-mail: gutwin@cs.usask.ca

Graphics Interface Conference 2007
28-30 May, Montréal, Canada

although there appears to be a relationship between the complexity of the visual data and the optimal transition time.

In the following sections, we review the ways that smooth transitions have previously been used in visualization systems, describe the two experiments, and discuss the implications of our results for designers of visualization systems.

2 RELATED WORK

Previous research on the use of animation for visual systems can be grouped into two general categories. From a bottom-up approach, some results report on the different ways that visual objects can be animated [1], on the use of artistic principles for designing appropriate animations [8], on the effectiveness of animated icons [2], or on the use of simple motion as a method for capturing attention [3]. From a top-down view, a number of studies have investigated the effectiveness of animation for teaching algorithms [16], for explaining complex concepts [9], or for understanding the cognitive benefits of animated displays in comparison to static representations [18]. While all these results can guide designers in producing better animated displays, they do not directly answer the question of whether smooth transitions assist users in working with visual information.

Below, we report on the visualization techniques that have used smooth transition in view changes, review studies that have inspired the work described in this paper, and report on the drawbacks inherent in smooth transitions.

2.1 Visualizations Benefiting from Smooth Transitions

A number of visualization systems have been developed using smooth transitions. However, the designers of these systems have used smooth animations to accomplish different objectives. Some objectives include making parts of the structure more visible, maintaining the perceptual relationships between different views, gradually increasing the visibility of the content, or collapsing and expanding visual structures.

Increasing structure visibility.

Several visualizations have used smooth animation to increase the visibility of structures. A classic example is the Cone Tree [11], a 3D representation of a hierarchy where the root of the tree is the apex of a cone and its children are evenly spaced around the circumference of the cone's base. This layout is iterated for the entire hierarchy. The 3D layout occludes nodes positioned further away from the user. As a result, the designers of the cone tree allow the user to see hidden structures by clicking on a node of interest. This smoothly rotates the tree in a period of less than a second to make the node and its path visible to the user.

Maintaining relationships between different views.

Zoomable user interfaces (ZUIs) have also explored the benefits of smooth transitions [6]. ZUIs allow users to zoom in to see details and to zoom out of a scene to see an overview. To allow the user to shift between multiple views, a number of ZUIs have integrated various types of animation in their interfaces. These animations help users understand where they are in the information space and in which direction they may be heading.

To maintain relationships between different views, smooth animations have also been employed in a number of focus+context visualization techniques. The general idea of using smooth transitions with focus+context systems is to facilitate a gradual shift in view between focus and context. One of the earliest focus+context visualization systems that uses smooth transitions is the perspective wall [11]. The perspective wall has three regions: a center region for viewing focused details and two perspective panels for viewing context. It provides smooth transitions to bring items of interest into the center region. Sunburst [17] is another focus+context visualization tool that uses smooth transitions to help the user maintain orientation during navigation. When users shift between detail and overview, the tool gradually shifts the view, assisting the user in identifying the part of the overview from which the details emerge.

The advantage of using smooth transitions between views is evident in Polyarchies [12], a complex visualization system that was designed to assist users in making sense of relationships that exist between multiple hierarchies. In a multi-part study, Robertson et al [12] compare various types of animations (sliding, horizontal rotations, stacked substrees) for showing the relationships between different structures. Their results show that a 'sliding' view that is based on smooth horizontal sliding of various hierarchical structures helps users maintain the visual relationships between the different views. Their results also suggest that animation speeds that complete a viewpoint change in one second are adequate for maintaining perceptual constancy.

Gradually increasing content visibility

Several techniques have used smooth transitions for gradually revealing information content. Continuous semantic zooming (CSZ) developed by Schaffer et al [13] is an example technique that employs smooth transitions to increase content visibility. This technique is characterized by two distinct but interrelated components: continuous zooming and presentations of semantic content at various stages of the zoom operation. When a region of interest becomes the focus, the user applies the continuous zoom to "open up" successive layers of the display. At each level of the operation the technique enhances continuity through smooth transitions between views and maintains location constraints to reduce the user's sense of spatial disorientation.

Continuous semantic zooming has been applied to information structures other than topological graphs. DateLens [5] employs CSZ to reveal varying degrees of content in tabular structures in a smooth and continuous manner. It applies linear distortions to cells of interest in a grid. As the level of distortion increases, semantic information is revealed based on the size of the region available for the display. An evaluation comparing DateLens to common calendar-based interactions reveals that continuous semantic zooming enhances content browsing in tabular structures [5]. Another distortion-based interactive technique was designed by Shi et al [15] for inspecting data in nodes of a TreeMap. The distortions are smooth transitions that gradually expand the space allotted to a node. This enables users to see elements at the leaf nodes without drilling-down through various layers of the hierarchy. In a study, Shi et al [15] showed that participants were able to identify content quicker and able to maintain context of the space better with smooth distortions.

Collapsing/expanding visual structures

A number of visualizations have benefited from smooth animations to expand information that was not previously visible or to collapse unnecessary structures that obstruct the view. Space Trees [10] is a hierarchical visualization system that combines the conventional layout of trees with a zooming environment that dynamically lays out branches of the tree to best fit the available screen space. Substructures of a tree that do not fit on the screen are summarized by a triangular preview. As the user clicks on the triangular preview, SpaceTree gradually expands the sub-structure and lays it out such that it takes maximum advantage of the screen space. In this technique, smooth transitions are used to aid the user in maintaining constancy between each level of the expansion/collapse of the substructure.

Elastic Hierarchies [21] are a hybrid visualization that combines treemaps with node-link diagrams. The design is motivated by using treemaps for their space-conserving properties and node-link subtrees for clarity in viewing the tree sub-structures. Smooth animation is employed in this visualization to expand a node-link view from a treemap view and to do the reverse. The authors suggest that using smooth transitions facilitates maintaining context when the visualization switches between different representation styles.

2.2 Animation for assisting the perception of connectivity and structural relationships

Only a limited number of studies have investigated whether smooth transitions assist users in maintaining connectivity and structural information in node-link diagrams. The goal of one experiment conducted by Ware and Bobrow [20] was to determine methods of assisting the perception of connectivity in large node-link diagrams. Different highlighting techniques were used to assist users in determining whether a pre-specified node was connected to a user-selected node using two major conditions. The static condition consisted of highlighting a certain number of edges away from the selected node. The animated condition consisted of setting into an oscillatory motion a subset of the graph that was connected to the selected node. Results showed that subjects performed equally well with either of the conditions, but performed significantly better when the motion and static cues were combined. In a second experiment, Ware and Bobrow [20] used the above two major conditions to determine whether subgraphs or paths within a larger graph intersected. The results showed that subjects could identify intersecting graphs if one was moving and the other was static. In both these experiments, motion was applied to parts of the graph to determine if substructures intersect or are connected. This form of motion is different than the use of smooth transitions that are of interest in this article. As such we cannot infer from their conclusions that smooth transitions between viewpoints will assist in maintaining perceptual connectivity or structural information.

Bederson and Boltman [4] conducted a study that had similar goals to the ones outlined in this paper. They examined how animating the change of viewpoint in a visual structure affects a user's ability to build a mental map of the information space. The authors compared two presentation types, animated and non-animated and designed three tasks for their experiment that tested the effectiveness of animation for forming spatial structures. For all their tasks, the participants were presented with a family tree containing images of different family members.

In the first task, subjects answered 9 questions about the relationships between family members by navigating through the family tree. Here participants were able to learn the relationships equally well with the animation as without. In the second task, participants navigated the family tree for 3 minutes and then answered 10 questions without looking at the family tree. This tested their ability to recall the information presented in the hierarchical structure. The authors hypothesized that if participants were able to build a better mental map with the animation, then they would be able to answer the questions more rapidly from memory. However, the results for this task also do not show any significant improvement in the smooth transition condition. Finally in the third task, the subjects assembled the structure of the family tree based on the contents of the nodes they had seen previously. In this task, subjects performed better with smooth transitions than without. However, the results showed an ordering effect, i.e., if smooth transitions were shown first, then they performed significantly better than if they were shown last.

The study by Bederson and Boltman [4] has several limitations and does not answer the questions we address here. First, they include semantic information within the structures (images of family members). Therefore reconstruction based on this information may shadow the effects observed with the different animation styles (Mary looks older, so she must be the mother or aunt of William). Second, the hierarchical structure that was used in the experiment was relatively small; the family tree contained four levels and nine nodes. Finally, the tasks in their study address whether users are able to formulate spatial information, but they do not deal with the issue of whether smooth transitions assist in maintaining connectivity and structural information in node-link diagrams.

2.3 Drawbacks of using smooth transitions

Smooth transitions have several drawbacks. The most significant drawback is that smooth transitions can take considerable amount of time to complete a viewpoint transformation, thereby increasing system response time [4]. The additional time taken may not benefit users who are familiar with the task or when the task is not complex. Another drawback with smooth transitions is that if they are not designed carefully, they can disrupt user performance and lead to distractions. In a series of experiments Bartram et al [3] evaluated the effectiveness of simple motion as a method of drawing the user's attention to an area of the display. Their results show that simple motion is significantly more effective than color or texture cues for distracting users. Their studies show that traveling motions which involve both detection and tracking are substantially more distracting than anchored motions. Their results also reveal that slow linear motion is irritating and distracting.

From a designer's perspective, smooth transitions also require more development effort. Additional algorithmic complexity is necessary to adequately interpolate between initial and final views of the animation. Furthermore, the designer may also need to consider details such as the display's refresh rate or the user's hardware capacity. These constraints put an additional overhead in the development effort required for building an animated system. In light of these drawbacks it is even more important for designers to be informed about the benefits that animations may provide. If there is evidence that animations provide significant benefits then designers may use these to outweigh the drawbacks of animated systems.

The limitations in prior studies, the apparent drawbacks of animations, along with the lack of strong empirical support for smooth transitions has motivated the work described here. The goals of this research are: 1) to quantify the effects of smooth transitions on perceptual constancy; 2) to determine the effect of transition speed in perceiving connectivity and recognizing structures; and 3) to design evaluation tasks that can adequately address questions 1 and 2 above. None of these goals have been addressed in previous research.

3 EXPERIMENT 1: CONNECTIVITY

Seeing connectivity in a graph or node link diagram is an essential perceptual task. In order to maintain perceptual constancy between views, it is important that the user be able to see and follow connections in a node-link diagram as the visualization undergoes smooth transitions. The objective of this experiment was to determine whether smooth transitions assist in perceiving connectivity in node-link diagrams. We predicted the following outcomes:

Hypothesis 1: users will be more accurate in perceiving connectivity when smooth transitions are applied to a viewpoint change of a node-link diagram.

Hypothesis 2: users will require less time to determine whether particular nodes in a node-link diagram are connected when smooth transitions are used.

3.1 Method

3.1.1 Subjects

Twelve subjects participated in this experiment (10 male, 2 female). Eleven subjects were computer science graduate students while one was an undergraduate student in the department of geography. All the participants were regular users of mouse- and windows-based systems and had four to ten years of experience with animated interfaces. All the subjects were exposed to animation though computer games. All the twelve subjects had seen or used planar graphs.

3.1.2 Materials

The graphs for this experiment were drawn in Microsoft Visio on a 900×900 pixel template with a white background. For experimental purposes, the template was divided into 9 cells, each with a size of 300×300 pixels, forming a 3×3 grid (Figure 2). Two types of graphs were used in this experiment: small and large. The small graphs were constructed with three nodes in each cell, while the large ones had six nodes in each cell. Each node was joined by a minimum of 3 and a maximum 6 links. Colors were used to differentiate the nodes in each grid. For the small graphs, three colors were used: red, blue and green, i.e., each cell contained only one red, only one blue and only one green node. Similarly, six colors were used for the larger graphs: red, light blue, dark blue, green, yellow and grey. In all the graphs, the links crossed over each other but did not cross over a node. Figure 2 below shows a small graph drawn on 3x3 grid. The grid lines were not shown to the participants during the actual experiment. In total there were 18 different small graphs and 18 different large graphs built using the criteria described above.

The experimental setup was developed using .NET running on a P4 Windows XP PC system. The display was a 17” monitor set to 1280×1024 resolution. The heart of this system was a viewport of size 300×300 pixels, showing one cell of the graph at any instance of time. Eight directional arrow buttons were provided for allowing the user to navigate through the entire graph. Clicking any one of the buttons would shift into the viewport another cell of the graph, corresponding to the direction indicated on the button, using either smooth or no transitions (details on video).

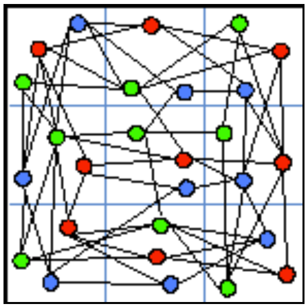

Figure 2 – Sample small size graph drawn on a 3×3 grid with all the red nodes connected.

3.1.3 Task

The task consisted of determining whether all the red nodes in the graph were connected in such a way that there always existed a link between two red nodes from adjacent cells. As the system showed only one cell of the graph through the viewport at any instance of time (Figure 3), the participants had to navigate through the entire graph using the eight directional buttons to see if all the red nodes were connected. If smooth transitions facilitate perceptual connectivity, then, in this condition, the participants should be able to determine very quickly and more accurately if two nodes, one of which is not in view, are connected. This task is representative of systems in which only part of the graph is visible at any instance. When smooth transitions were employed, the participant was able to see the current cell of the graph move smoothly out of the viewport and the next cell move smoothly into the viewport. This scenario is pictorially depicted in the Figure 3 below. Figure 3.a shows the initial cell of the graph through the viewport and Figure 3.d shows the final cell of the graph when the graph is shifted to the left. Figures 3.b and 3.c show a snapshot of the contents of the viewport during transition.

In contrast, when no transitions were employed, the participant would not see the intermediate views of the graph. The net effect was that, in the no-transition mode, the users only saw the views in Figures 3.a and 3.d. Clicking on the directional arrow buttons moved the viewport to the next cell of the graph in the corresponding direction, thereby showing a different subgraph. The task of the participant was to follow the links coming out of a red node and to determine if any of the links are connected to another red node in a different cell. The participant was allowed to navigate through the graph any number of times, until he/she was comfortable answering whether all the red nodes were connected. They answered this by clicking on either of the two buttons (‘YES’ and ‘NO’) that were provided near the arrow buttons. We also collected the number of moves they required to answer the question and used the total number of moves, to determine which method took them the longest.

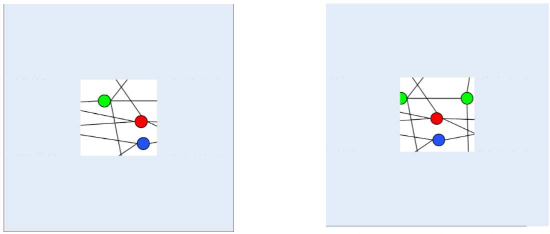

(a) initial view **(b) intermediate view #1**

(c) intermediate view #2 **(d) final view**

Figure 3 – Snapshot of graph in Figure 2 moving to the left.

3.1.4 Design

The experiment was setup using a 3×2 within-participants factorial design. The factors are:

- Transition style: Slow-Transition, Fast-Transition, No-transition. In prior literature researchers have hinted at 0.5 secs being the lower limit on our ability to maintain perceptual constancy [4,11]. We used this metric for the speed rate of a fast transition. Prior results also suggest that an upper limit of 1 sec [4] can be disruptive and lead to impaired performance. We doubled this limit and selected a rate of 2 secs for the slow transition rate. The conditions are:
    - *Slow-Transition*: this style used an animation speed of 150 pixels per second. This corresponds to a movement of 2 secs to refresh the viewport.
    - *Fast-Transition:* this style used an animation speed of 600 pixels per second. This corresponds to a movement of 0.5 secs to refresh the viewport.
    - *No-Transition:* no animation is used.
- Graph size: Small (3 nodes per cell or 27 nodes), Large (6 nodes per cell or 54 nodes)

Transition style was fully counterbalanced using a Latin square design. The other factor was always presented in increasing order (i.e., from smaller to larger graphs). Within each condition, participants carried out 6 trials. With 12 participants, 3 transition styles, 2 graph sizes and 6 trials per condition, the system recorded a total of 432 trials. The system collected the total number of moves through the graph, the errors and the completion time. The completion time, is the time from the moment the participant starts the task of navigating the entire structure to the point of responding YES/NO by clicking on the buttons. Participants also filled out a brief questionnaire regarding their preferences at the end of the experiment.

3.1.5 Procedure

Participants were randomly assigned to one of the six order groups obtained by counterbalancing the transition styles. Prior to starting the experiment, participants were given a small practice session which involved 2 trials per condition. After completing the practice trials, all participants indicated that they were comfortable with the three transition styles. The participants then completed 36 trials without any breaks. At the end of the trials, the participants were asked to indicate the transition style that was easiest and the style for which they felt they performed the fastest.

3.2 Results and Discussion

To test the hypotheses stated in the beginning of this section, we measured subjects' performance on the given task with respect to errors and number of moves.

Error Rate

The average error rate is summarized in Figure 4 below. The error rate was analyzed by means of a two-way ANOVA, with both Transition Style (Slow-Transition, Fast-Transition, No-Transition) and Graph Size (Small and Large) serving as repeated measures. An alpha level of .05 was used for all statistical tests. Pair-wise comparisons were performed using Tukey HSD tests. Transition style was found to be significant (F(2, 22) = 42.39, p < .001) with the fast-transition average error rate (9%) being smaller in comparison to the slow-transition (10.4%) and the no-transition (35.4%) error rates. The main effect for Graph Size was significant (F(1,11) = 0.096, p = .763). The interaction effect was not significant (F(2,22) = 2.2, p = .135). Pair-wise comparisons show that there is a significant difference between slow transitions and no transitions (p < .001) and between fast transitions and no transitions (p < .001). However, there is no significant difference between fast transitions and slow transition (p < .689).

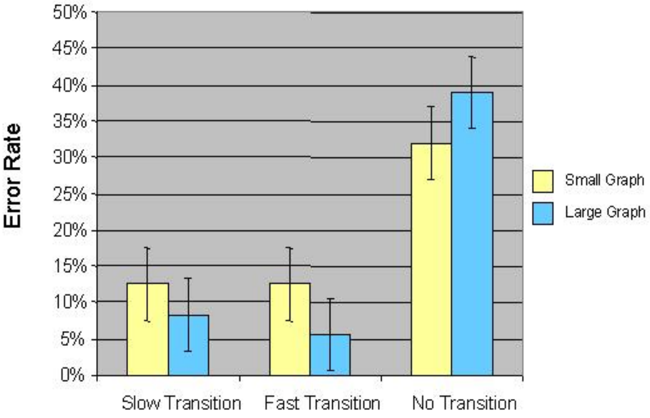

Figure 4 – Average error rates (with s.e.d.) by transition style and graph size.

We observe that the participants made ¼ of the errors in smooth transitions as they did without transitions. The results support the first hypothesis in that participants will be able to perceive connectivity more accurately with smooth transitions.

Completion Time

The average completion time for each condition is summarized in Figure 5. Completion time was analyzed by means of a two-way ANOVA, with both Transition Style (Slow-Transition, Fast-Transition, No-Transition) and Graph Size (Small and Large) serving as repeated measures. Transition style was found to be significant (F(2, 22) = 4.751, p = .019) with the participants completing the task faster with the fast-transition condition (37.4 secs) than the slow-transition condition (49 secs) or the no-transition condition (50.9 secs). The main effect for Graph Size was significant (F(1,11) = 8.272, p = .015). We did not observe any interaction effects (F(2,22) = 1.22, p = .314). Pair-wise comparisons show that there is a significant difference between slow-transitions and fast-transitions (p = 0.002) and between fast-transitions and no-transitions (p = .044). However, there is no significant difference between slow-transitions and no-transition (p = .707).

Figure 5 – Completion times (with s.e.d.) by transition style and graph size.

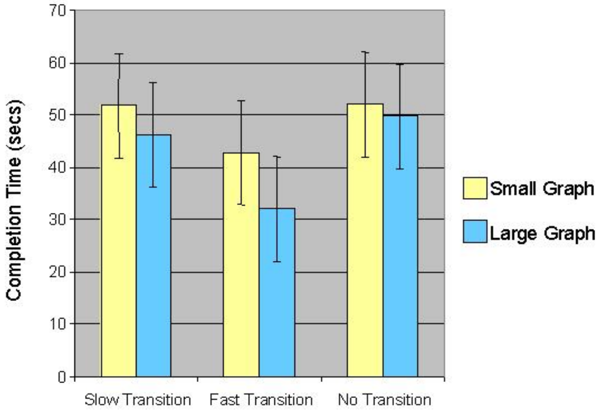

The analysis on completion time provides some very strong implications. In particular, the results suggest that it can take users less time to process connectivity information with smooth animations. These results are remarkable considering that for each transition, the system response time is significantly higher with smooth transition than without. By comparing error rates and completion times, we observe that faster transitions (0.5 secs) provide further gains than slower transitions; the error rates are approximately the same while completion times are lower with faster transitions. However, these observations may be directly linked to the complexity of the task, i.e. a more complex task may require slower transitions. In the second experiment we investigate the effects of transitions under a more complex task.

4 EXPERIMENT 2: STRUCTURE RECOGNITION

The first experiment was designed to determine whether connectivity is perceived more easily with smooth transitions than without. Another method of measuring the effectiveness of smooth transitions is to see whether they assist in forming structural relationships. Experiment 2 was designed to investigate the effects of smooth transitions in recognizing structures. The study was designed to determine whether the "whole" can be constructed from its visible "parts". The task was inspired by Biederman's design for testing recognition-by-component theory (RBC) [7]. Biederman's results show that the human perceptual system is capable of recognizing objects by simply identifying a few major components of the object's structure, i.e., the structure can be reassembled from its parts [7]. In this way, recreating the whole from its parts can assess the effectiveness of a modality for assisting in recognizing structures [19]. This task requires a more significant cognitive effort than the task in experiment 1. We predicted the following outcomes:

Hypothesis 1: participants will reconstruct the structure more accurately with smooth transitions.

Hypothesis 2: participants will reconstruct the structure more rapidly using smooth transitions.

Hypothesis 3: participants will perform better (speed and accuracy) with slower transitions than with faster transitions.

Hypothesis 4: participants will more easily (speed and accuracy) recognize simpler structures than complex structures throughout all conditions.

4.1 Method

4.1.1 Subjects

Twelve paid volunteers (7 male, 5 female) participated in this experiment. All the participants were recruited from a local university. Eight of them were graduate students in Computer Science, two were graduate students in Mathematics, and the remaining two were undergraduates from the Faculty of Arts. Though all were regular users of mouse- and windows-based systems (at least four hours per day), their experience in using animated interfaces varied from four to fifteen years. All of them had experience with animation primarily through computer games. All the participants were also familiar with trees.

4.1.2 Materials

Two types of trees were used for this experiment: shallow and deep. The shallow trees were constructed using three levels (the root node being at level 0) and the deep ones using five levels. All the trees were drawn using Microsoft Visio on a 900×900-pixel template with a white background. For experimental purposes the template was divided into 9 cells, each of size 300×300 pixels, thereby forming a 3×3 grid. The nodes and links in the tree were drawn in black. Figure 6 shows a sample 5-level tree (deep tree) drawn on a 3×3 grid (Figure 6). The grid lines were not shown to the participants during the actual experiment.

The experimental system was similar to that of experiment 1. In this experiment participants were give one button (the MOVE button) to navigate around the graphs. When the user clicked on MOVE the system shifted into the viewport another cell of the tree either with or without transitions.

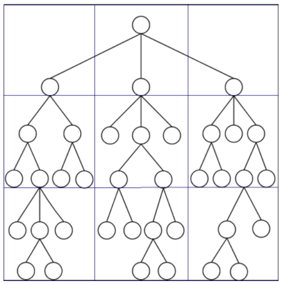

Figure 6 – Sample deep hierarchy used in experiment 2.

4.1.3 Task

The system, at any instance of time showed only one substructure of the tree to the participant through the viewport. Clicking on the MOVE button would shift the graph such that a randomly selected portion of the tree entered the viewport. The task of the user was to remember each substructure to the best of their ability and knowledge so that they could recognize the entire structure after seeing all the parts of the tree. When using smooth transitions, the participant was able to see the initial substructure move smoothly out of the viewport, and the next substructure move smoothly into the viewport. This scenario is similar to the one shown in experiment 1 with the only difference that in this case a subtree is moving out and replaced by another subtree of the hierarchy. The user continued to press the MOVE button until the system automatically stopped the transitions after presenting the entire tree twice. After seeing all the parts of the tree through the viewport, the user was presented with four trees and was asked to select the tree structure that was composed of the subtrees seen through the viewport.

The four trees were carefully constructed so that they differed in the following manner:

- Tree 1 - original: This tree is the same as the original tree shown in the viewport (Figure 7.a). The user would have to select this tree to get a correct score.
- Tree 2 – one-node difference: This tree differs from the original tree in only one node (Figure 7.b). One node is either removed or added to the original tree to get a one-node difference tree.
- Tree 3 – two-node difference: A two-node difference tree differs in two nodes as compared to the original tree (Figure 7.c). Either two nodes are removed or two nodes are added to the original tree to get this type of a tree.
- Tree 4 – layout difference: A layout difference is formed by interchanging the positions of two subtrees in the original tree (Figure 7.d).

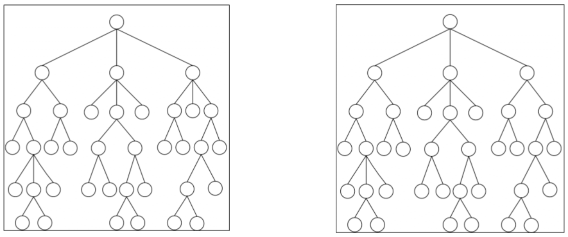

(a) Original tree **(b) One-node difference**

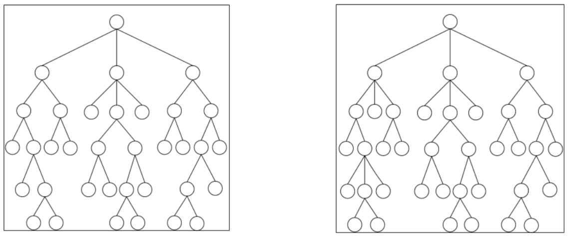

(c) Two-node difference **(d) Layout difference**

Figure 7 – Four choices for the tree in Figure 6.

4.1.4 Design

The experiment used a 3x2 within-participants factorial design. The factors consisted of *transition style* and *tree size*.

- Transition style: the same styles, slow-transition, fast-transition and no-transition, as in experiment 1 were used.
- Tree size: s*hallow* - consisting of a hierarchy with 3-levels, d*eep* - consisting of a hierarchy with 5-levels.

Transition style was fully counterbalanced using a Latin square design. Tree size was always presented in increasing order (i.e., from shallow to deep trees). Within each condition, participants carried out 6 trials. With 12 participants, 3 transition styles, 2 tree sizes and 6 trials per condition, the system recorded a total of 432 trials. The system collected accuracy, the type of error if the participant chose the wrong tree and response times. The response time is defined as the time from the moment the participant is shown the four tree structure to the moment the participant enters his/her response. Participants also filled out a brief questionnaire regarding their preferences at the end of the trials.

4.1.5 Procedure

Participants were randomly assigned to one of the six conditions (transition style × tree size). The procedure used was similar to the one used in experiment 1.

4.2 Results and discussion

To test the three hypotheses stated in the beginning of this section, we measured subjects' performance on the given task with respect to errors and response time.

Error Rate

The average error rate is summarized in Figure 8. The error rate was analyzed by means of a two-way ANOVA, with both Transition Style (Slow-Transition, Fast-Transition, No-Transition) and Tree Size (Shallow vs. Deep) serving as repeated measures. Pair-wise comparisons were performed using Tukey HSD tests.

Transition style was significant (F(2, 22) = 25.05, p < .001) with the slow-transition mean error rate (26.4%) being smaller than the fast-transition (30.6%) and the no-transition (66.7%) error rates. The main effect for Tree Size was not statistically significant at the 0.05 level (F(1, 11) = 4.475, p = .058). However, a significant interaction effect was found between Transition Style and Tree Size (F(2, 22) = 9.725, p = .001).

Pair-wise comparisons reveal that the error rate with slow-transition is significantly lower than the error rate with no-transition (p < .001). Similarly, the error rate with fast-transition is significantly lower than the error rate with no-transition (p < .001). The results support hypothesis 1 in that participants are more accurate in reconstructing the structure with smooth transitions (fast or slow) than without. However, the error rate with slow-transition is not significantly lower than the error rate with fast-transition (p=.236). This suggests that transition speed had no effect on accuracy, thereby not supporting hypothesis 3.

Participants are more accurate in reconstructing the 3-level tree than the 5-level tree (p = .044). This supports hypothesis 4; participants perform more accurately when the structure is smaller. Interestingly, we notice in the no-transition condition that participants are less accurate with the smaller structure than the larger structure (but no significance is found). One reason for this could be due to learning effects as the deeper trees were presented after the participants completed the trials with the shallow trees.

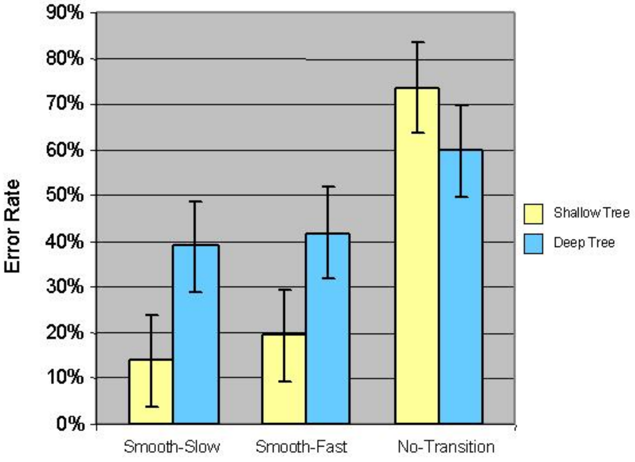

Figure 8 – Average error rate (with s.e.d.) by transition style and tree depth for the task of reconstructing the hierarchy.

Response Time

The average response time is summarized in Figure 9. The response time was analyzed by means of a two-way ANOVA, with both Transition Style and Tree Size serving as repeated measures. An alpha level of .05 was used for all statistical tests. Pair-wise comparisons were performed using Tukey HSD tests.

The main effect for transition style was not found to be significant (F(2, 22) = .109, p = .898). The main effect for tree size was statistically significant, (F(1, 11) = 11.259, p = .006) with the average response time for the 3-level tree at 16.25 secs and for the 5-level tree at 23.324 secs; this supports hypothesis 4. A significant interaction effect was not found between transition style and tree size, F(2, 22) = 3.317, p = .055. These results do not support hypothesis 2 in that participants are not faster in responding with smooth transitions than without. Interestingly, we observe that on average participants took longer to respond to the 5-level trees with smooth transitions than without. Pair-wise comparisons show that there is no significant difference in response time between slow-transitions and fast-transitions (p=.254). This does not support hypothesis 3 in that transition speed has no effect on performance.

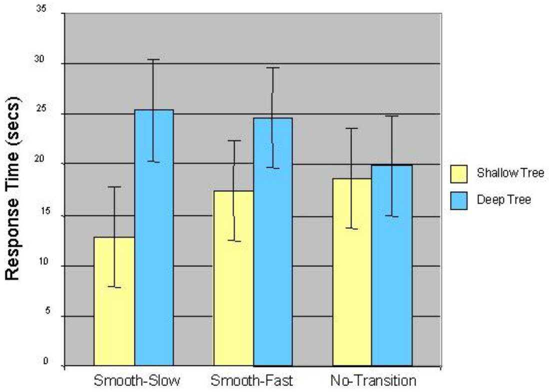

Figure 9 – Average response time (with s.e.d.) for selecting the correct type of structure.

Overall, the analysis of error rates and response times suggest the following: a) participants are able to reconstruct the structure more accurately with smooth transitions than without (support hypothesis 1); b) smooth transitions do not have an effect on how rapidly participants are able to reconstruct the structure (reject hypothesis 2); c) participants do not perform better with slower transitions than with faster transitions (reject hypothesis 3); d) participants are able to reconstruct simpler structures more easily than complex structures (supports hypothesis 4).

5 DISCUSSION

The results provide evidence that smooth transitions have a strong effect on perceptual constancy. The findings show that:

- Smooth transitions assist in maintaining connectivity between different views of a node-link diagram,
- Smooth transitions enforce the gestalt of continuity thereby reducing the amount of cognitive effort required to reformulate the "whole" from the "parts",
- Performance does not differ when views of node-link structures undergo, within reasonable limits, slow or fast transitions. This means that even with faster transitions (of upto 0.5 secs for a view change) users are able to maintain perceptual constancy.

In the following sections we discuss the reasons for these results and summarize some main guidelines for designers.

Reasons for findings

The first experiment focused on the use of smooth transitions to establish connectivity in node-link diagrams. As the participants were asked to find if all the red nodes were connected, they concentrated only on the links coming out of or moving into a red node. Participants felt this was an easy task when smooth transitions were employed. The main strategy in this task consisted of visually following the links coming out of a red node and seeing which edge connects to another red node in the next cell of the graph. With smooth transitions they were capable of following the edges over the transition period. In contrast, without smooth transitions, the participants had to move back and forth between the two views to confirm the connectivity between two nodes. This resulted in longer completion times and more errors. The second experiment relied on a participant's cognitive ability to reconstruct a tree after seeing subtress of the hierarchy. The smooth transitions helped the participants to easily remember the parts of the tree and the relative positioning of these parts in the entire tree structure. Although, in the no-transition condition, they remembered the various subtrees, they were not able to reconstruct the entire tree. This is mainly attributed to the lack of orientation information the users possessed without transition.

Interestingly, our results show that participants took an equal amount of time between viewpoint changes with smooth transitions as without transitions. On average the participants were required to bring parts of the node-link diagram into view more frequently without smooth transition so that they could formulate the structures and maintain perceptual constancy. This suggests that longer delays with smooth transitions do not really constitute a tradeoff but instead leverage the user's cognitive resources that would otherwise be used in reorganizing the structures.

In both experiments, over 80% of the participants preferred the smooth transition conditions over no-transitions. However, we noticed that in the first experiment three times as many participants preferred the fast transition over the slow transition. In the second experiment more participants preferred the slow transition condition to the fast transition condition. Since the task in the second experiment is more complex, participants required more time to assimilate the various views. Although this requires further investigation, we believe that there exists a strong correlation between the complexity of the task and the user's preference of transition speed.

Guidelines for designers

We believe that there are some guidelines designers of visual workspaces might find valuable from our findings:

- Use smooth animated transitions for facilitating perceptual constancy for viewpoint changes of node-link diagrams.
- Slow transitions do not provide significantly more benefits than fast transitions. Therefore for simple tasks, faster transitions should be used as they reduce delays and provide similar benefits as slower transitions.
- For complex tasks, slower transitions are preferred to fast transitions. Therefore, interfaces that require explicit manipulation of viewpoint changes should give users control over the transition speed. The specific relationship between task complexity and animation time will be explored in future work.

6 CONCLUSION

Smooth transitions are very common in many visual systems. Designers are primarily guided and motivated by intuition when they use smooth animations. In this paper we describe two experiments that were designed to quantify the effect of smooth transitions on perceptual constancy in node-link diagrams. Our results show that participants are able to perceive connectivity better when transformations of node-link diagrams undergo smooth transitions than when they do not. Our results also show that smooth transitions assist in maintaining structural information in viewpoint changes of node-link diagrams. An important contribution of this work is that viewpoint changes of node-link diagrams that take place in 0.5 seconds facilitate perceptual constancy. This result is important as it can guide designers in integrating smooth transitions in visual systems. In future work, we plan to investigate the use of smooth transitions in establishing connectivity and structural formation in 3-dimensional node-link diagrams. We also intend on quantifying the effects of smooth transitions with zooming or other interactive tasks, on determining the correlation between transition speed and task complexity, and in investigating the effects of different transition styles, such as slow-in/slow-out or variable transitions speeds on task performance.

REFERENCES

[1] R. Baecker and I. Small. Animation at the interface. In B. Laurel, editor, *The Art of Human-Computer Interface Design*, Addison-Wesley, MA, pp. 251-267, 1990.

[2] R. Baecker, I. Small, & R. Mander. Bringing Icons to Life. *ACM CHI*, pp. 1-6, 1991.

[3] L. Bartram, C. Ware, W.T. Calvert. Moticons: detection, distraction and task. *Int. J. Hum.-Comput. Stud.* 58(5): 515-545, 2003.

[4] B.B. Bederson, A. Boltman. Does animation help users build mental maps of spatial information?, *IEEE InfoVis*, pp. 28 – 35, 1999.

[5] B. B. Bederson, A. Clamage, M. Czerwinski, G. G. Robertson. DateLens: A fisheye calendar interface for PDAs, *ACM TOCHI*, 11(1):90-119, 2004.

[6] B. B. Bederson, J. D. Hollan, K. Perlin, J. Meyer, D. Bacon, G.W. Furnas. Pad++: A Zoomable Graphical Sketchpad for Exploring Alternate Interface Physics. *Journal of Visual Languages and Computing*, 7(1): 3-32, 1996.

[7] I. Biederman. Recognition-by-Components: A Theory of Human Image Understanding. *Psychological Review*, 94(2): 115-147, 1987.

[8] B. W. Chang, D. Ungar, Animation: From Cartoons to the User Interface. *ACM UIST*, pp. 45-55, 1993.

[9] C. Gonzalez. Does Animation in User Interfaces Improve Decision Making? *ACM CHI*, pp. 27-34, 1996.

[10] C. Plaisant, J. Grosjean, B. B. Bederson. SpaceTree: Supporting Exploration in Large Node Link Tree, Design Evolution and Empirical Evaluation, *IEEE InfoVis*, pp. 57-64, 2000.

[11] G. G. Robertson, S. K. Card, J. D. Mackinlay: Information Visualization Using 3D Interactive Animation. *CACM*, 36(4): 56-71, 1993.

[12] G. G. Robertson, K. Cameron, M. Czerwinski, D. C. Robbins. Polyarchy visualization: visualizing multiple intersecting hierarchies. *ACM CHI*, pp. 423-430, 2002.

[13] D. Schaffer, Z. Zuo, S. Greenberg, L. Bartram, J. Dill, S. Dubs, M. Roseman. Navigating Hierarchically Clustered Networks Through Fisheye and Full-Zoom Methods. *ACM TOCHI* 3(2):162-188, 1996.

[14] A. Sears, J.A. Jacko, & M.S. Borella, Internet Delay Effects: How Users Perceive Quality, Organization, and Ease of Use of Information, *Extended Abstracts of CHI*, pp. 353-354, 1997.

[15] K. Shi, P. Irani, B.Li. An Evaluation of Content Browsing Techniques for Hierarchical Space-Filling Visualizations.*IEEE InfoVis*, pp. 81-88, 2005.

[16] J. T. Stasko, A. Badre, & C. Lewis. Do Algorithm Animations Assist Learning? An Empirical Study and Analysis. *ACM CHI*, pp. 61-66, 1993.

[17] J. T. Stasko, E. Zhang. Focus+Context Display and Navigation Techniques for Enhancing Radial, Space-Filling Hierarchy Visualizations. *IEEE InfoVis*, pp. 57-65, 2000.

[18] B. Tversky, J.B. Morrison, and M. Betrancourt, Animation: can it facilitate? *IJHCS*, 57(4): p. 247-262, 2002.

[19] C. Ware, *Information Visualization: Perception for Design*. Morgan Kaufmann Publishers, San Francisco, 2004.

[20] C. Ware, R. J. Bobrow. Supporting visual queries on medium-sized node-link diagrams. *Information Visualization* 4(1): 49-58, 2005.

[21] S. Zhao, M. J. McGuffin, M. H. Chignell. Elastic Hierarchies: Combining Treemaps and Node-Link Diagrams. *IEEE InfoVis*, pp. 57-64, 2005.