Bimanual and Unimanual Image Alignment: An Evaluation of Mouse-Based Techniques

Celine Latulipe *Craig S. Kaplan* *Charles L.A. Clarke*
Computer Graphics Laboratory
University of Waterloo
{clatulip, cskaplan, claclark}@uwaterloo.ca

ABSTRACT

We present an evaluation of three mouse-based techniques for aligning digital images. We investigate the physical image alignment task and discuss the implications for interacting with virtual images. In a formal evaluation we show that a symmetric bimanual technique outperforms an asymmetric bimanual technique which in turn outperforms a unimanual technique. We show that even after mode switching times are removed, the symmetric technique outperforms the single mouse technique. Subjects also exhibited more parallel interaction using the symmetric technique than when using the asymmetric technique.

ACM Classification H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General Terms Design, Human Factors

Keywords: Bimanual input, two-handed interaction, symmetric interaction, image registration, image alignment

INTRODUCTION

Bimanual interaction has been studied extensively and many bimanual interfaces have been proposed since the first two-handed study by Buxton and Myers [3]. More recently, MacKenzie and Guiard have pointed to the lack of affordances for bimanual interaction provided by current desktop interfaces [19]. The use of specialized devices (tablets, datagloves, flocking devices, trackballs or six degree-of-freedom mice) in the previously studied bimanual interfaces is partly responsible for preventing the widespread adoption of bimanual input. We believe that the use of inexpensive, standard devices will be instrumental in bringing two-handed input to the everyday computer desktop. In this paper we propose and evaluate the use of two standard USB mice and a symmetric two-handed interaction technique for tasks that require basic geometric transformations of two-dimensional objects. Figure 1 shows our implementation of this technique in our two-handed drawing program, symDraw [16]. The two cursors are attached to opposite corners of the object and by moving the cursors around the screen, the object is simultaneously scaled, rotatated and translated.

Although users may not think in geometric terms, they are performing geometric transformations when they drag a file folder, resize a window, select some text or rotate an image. In current graphical user interfaces, user tasks reflect the affordances of a single spatial input device. With a single mouse or trackball, a user can perform either a translation, a scale or a rotation, but cannot easily combine any of these actions. A scale and rotation can be combined only if they both occur with respect to some previously selected fixed point.

Experienced computer users have become accustomed to the paradigm of performing simple geometric transformations in serial. However, this is unnatural compared to the way people combine translation, rotation and scaling seamlessly in everyday interactions with objects. An elastic band, for example, can be rotated, translated and scaled simultaneously. If you want to put an elastic band around an object, you use your hands to stretch the elastic band, while translating and rotating your hands together to get it in place around the object. Even when objects cannot be scaled, we effortlessly translate and rotate them at the same time. It is awkward, and sometimes impossible, to serialize these manipulations. And yet, computer users routinely manipulate on-screen objects this way.

The addition of a second spatial input device alleviates the need to switch between geometric transformations and permits transformations that combine scale, rotation and translation. The main question that arises is how to map the two degrees of freedom from each spatial device to the three geometric transformations. An asymmetric mapping might use the non-dominant hand to control the translation of an object, while the dominant hand controls scale and rotation. A symmetric mapping would involve the two hands controlling opposite corners of the object to 'steer' it through space, achieving simultaneous scale, translation and rotation. We believe that the symmetric mapping is likely to yield the highest performance for these tasks.

This paper aims to evaluate both symmetric and asymmetric dual-mouse techniques against the typical one-handed technique with mode-switching. We have chosen the task of aligning digital images for this evaluation. Pasting together separate digital photographs into a panoramic image in a typical graphics program is difficult because the images often

UIST'05, October 23-27, 2005, Seattle, Washington, USA.

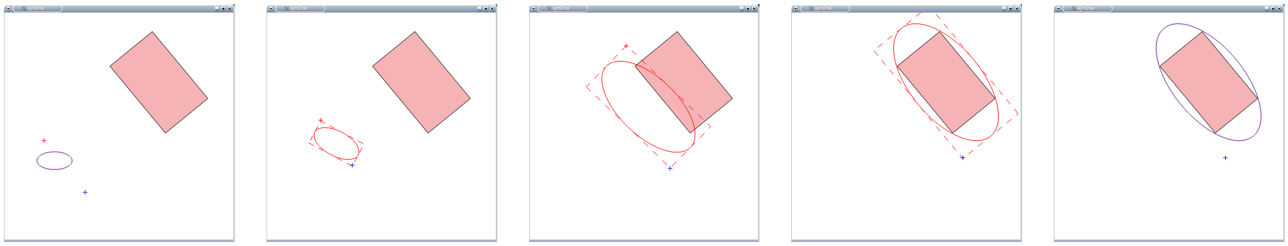

Figure 1: Symmetric interaction allows objects to be scaled, translated and rotated simultaneously, simplifying tasks such as enclosing a rectangle within an oval.

need to be slightly scaled, slightly rotated and slightly overlapped. We believe the manual image alignment task is an appropriate choice for our evaluation because it is a task which should be easy for subjects to understand, and so we expect the learning effects to be minimal.

In this paper, we first review previous work in the areas of two-handed interaction, both asymmetric and symmetric. We then discuss the image alignment task and examine the differences between physical and virtual image alignment and the implications for interaction. We then describe an experiment to test the effectiveness of mouse-based unimanual and bimanual techniques in the image alignment task. We discuss the results of this experiment and give some conclusions and directions for future work.

PREVIOUS WORK

Research in the area of two-handed input has generally focused on asymmetric interaction techniques. Examples include Shaw's THRED system [22] and Fitzmaurice's graspable bricks [10], as well as many others [6, 7, 12, 13, 21]. Design and evaluation of such techniques have most often been guided by the principles set out by Guiard [11]. Guiard characterizes asymmetric interactions as those in which the non-dominant hand (NDH) sets the frame of reference in which the dominant hand (DH) works. Additionally, the action of the NDH precedes that of the DH. The level of detail, both spatial and temporal, at which the DH works is finer than that of the NDH. Guiard's work is important because many everday activities, such as handwriting and sweeping, can be characterized as asymmetric.

However, there are some activities that people perform regularly that cannot be characterized as asymmetric. Skipping rope is an activity in which both hands work together at the same level of spatial and temporal detail. Both hands start and end together. Other examples of symmetric interaction include typing at a keyboard and folding linens.

Despite the fact that symmetric interactions exist in real life, there has been little research into the use of symmetric interaction techniques in the user interface. One notable example of research that takes advantage of symmetric interaction is the deformation technique found in the Twister implementation for 3-D object manipulation [18]. In their system, the user can pull and twist an object using both hands. However, the symmetry of the interaction was not explored in that research and the implementation used specialized input devices. Casalta et al. compared symmetric and asymmetric techniques for drawing rectangles and found better performance with the symmetric technique [5]. They questioned the validity of Guiard's kinematic chain model when the two hands are working on the same object. Owen et al. [20] studied bimanual and unimanual curve manipulation using a digitizing tablet with a puck in the NDH and a stylus in the DH. They found that the two-handed technique (which involved symmetric movements with asymmetric devices) was faster than the single-mouse technique. However, their technique involved the two hands manipulating separate objects (different spline control points) and their bimanual results were worse than the unimanual results once the switching costs were removed. The results from these studies suggest that symmetric interaction is best suited to manipulation of a single object.

Balakrishnan and Hinckley studied a symmetric target following task in which users had to track two different targets around the screen using two cursors [1]. They measured both the level of parallelism and the level of symmetric interaction exhibited by subjects. They found that reducing the amount of visual integration between the two tracked objects led to degradation both in the symmetry of action and in the level of parallelism. They concluded that objects being manipulated in a symmetric interaction should be visually connected. In image alignment, this recommendation is met if the two hands are manipulating the same image.

The idea of using the opposite corners of an object to scale, rotate and position the object was briefly described in the T3 system [15], but was not quantitatively evaluated. A similar symmetric interaction technique using only scaling and translation (but not rotation) was evaluated for an area sweeping task [17]. That research articulated two sources of benefit of two-handed interaction over one-handed interaction: the manual benefit of moving two hands in parallel, and the cognitive benefit of a unified task representation.

In summary, symmetric interaction techniques have not been studied as extensively as asymmetric techniques, but the few studies that have examined symmetric interaction have reported positive results. Our research aims to augment this work and evaluate the effectiveness of symmetric interaction for tasks that require rotation, translation and scale.

IMAGE ALIGNMENT

The task we have chosen to use in our evaluation is image registration. Image registration is commonly used in photography, digital art, medical imaging and remote sensing. In image registration, two or more images must be aligned to form a single complete image. In certain domains, such as remote sensing, there is no user interaction, and the images are processed by specialized software that expects a certain ordering and specific camera angles. However, for photographers and artists, the alignment of images is often done manually. In some software, algorithms are used after manual alignment to perfect the registration and adjust factors such as lighting and shadows. We choose to concentrate on evaluating the interactive aspect of image alignment.

There is little published research in the area of manual image alignment; the research in this area has focused on automatic image registration algorithms [8, 23]. Instead, one must look to commercially available software to see what manual techniques are used. These techniques fall into two categories: manual movement of the images on a virtual desktop (such as in Photoshop and PTgui [4]), and point-registration methods (such as in Panavue [9]), where the user selects coincident points on the images to be aligned. We focus on the image registration techniques involving manual manipulation of images though translation, rotation and scale.

Physical Versus Virtual Image Alignment

It is worth examining how the manual image alignment task is done outside of the virtual workspace. We ran an informal experiment, in which users were asked to perform the task with different types of paper. Small, light-weight images were translated and rotated by one hand. For larger and/or heavier images, we observed that the image was roughly translated into place with each hand holding an edge or corner and then simultaneously translated and rotated into final position using both hands symmetrically. There are two important points to be gleaned from this experiment. When both hands are used, they are used symmetrically. And, regardless of the size of the picture, the fine-tuning part of the interaction involved *simultaneous* rotation and translation.

A single hand can simultaneously rotate and translate a small object such as a photograph in three dimensions. In the virtual task, a single mouse provides only two degrees of freedom for translation in a two-dimensional space. Thus, simultaneous rotation and translation of the image is theoretically not possible using only one mouse, though some pseudo-physics algorithms have been proposed [2, 14]. Specialized devices such as digital pucks could provide affordances similar to the hand-wrist combination, but we choose to investigate what can be accomplished with standard mice.

The second interesting difference between the physical alignment task and the corresponding virtual task is that the virtual task allows the scaling of images, which is not possible in the physical realm. If we include the ability to scale images in the virtual image alignment task, which seems desirable, the single-mouse pseudo-physics algorithms for simultaneous translation and rotation are insufficient. The second mouse is necessary to achieve the desired effect.

Physical image alignment demonstrates that virtual image alignment is suited to symmetric interaction. Additionally, scaling is a task that is naturally symmetric: the distance between the cursors can be used to control scale. Thus we expect the symmetric image alignment technique, where the two cursors are attached to opposite corners of the image, to outperform single-mouse techniques and dual-mouse asymmetric techniques.

EXPERIMENT

We designed an experiment to test symmetric, asymmetric and single-mouse interaction for the image alignment task.

Experimental Design

The experiment was conducted as a within-subjects comparison of three conditions: single-mouse, dual-mouse asymmetric and dual-mouse symmetric. Twenty-four subjects, all undergraduate students not majoring in computer science or computer engineering, were paid $10 to complete the experiment. 13 subjects were male, 11 were female and all self-declared as right-handed. For each condition, four images were multiplied by ten transformation configurations to produce 40 different trials which were presented in random order. Each subject completed four practice trials plus 40 measured trials in each of the three conditions, for a total of 132 trials. The three conditions allowed six different presentation orders, and the subjects were grouped and counterbalanced so that four subjects saw each presentation order.

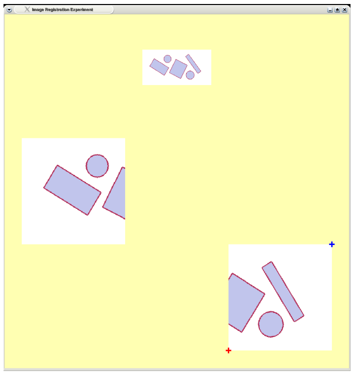

Figure 2: At the start of each trial the two images to be aligned are displayed, as well as a small complete image at the top of the screen, so that subjects know what the final picture should look like.

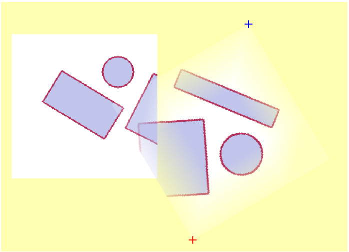

Figure 3: As Image B (right) is manipulated, its edges become transparent to help in the overlapping aspect of the image alignment task.

Task

In each experiment trial, the subject was asked to align two images into a single, complete picture. Each trial began with the appearance of one or two small start circles. The subjects had to move their cursor(s) into the start circles. This was followed by four 'readiness' beeps. Then the images would appear and the trial timing would start (see Figure 2). Image A was always on the left side of the screen, and could not be moved. Image B was initially located in the bottom right of the screen. The subjects were required to translate, rotate and scale Image B until it was aligned with Image A. Image A displayed exactly one half of the source image and was not rotated or scaled. Image B was jittered according to one of ten transformation configurations described later in this section. Image C, a small version of the full source image, was displayed at the top center of the screen, to give subjects an idea of how the aligned images should look. To aid in the registration process, the edges of Image B were made transparent while it was being manipulated. This allowed subjects to see both Image B and part of Image A underneath it (see Figure 3). When the subject had aligned Image B to Image A, a red border appeared around Image A to signify that the alignment was complete (see Figure 8). After 200 milliseconds, the images disappeared and the trial was over.

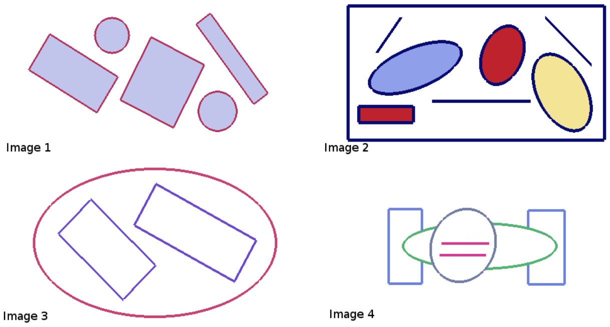

Figure 4: The four graphics images that were used in the experiment. Each graphic was treated as an OpenGL texture and transformed onto two rectangles to create image pairs that subjects were asked to align.

Image Sets Four different source images were used for generating image sets and each image was seen 33 times by each subject: three conditions multiplied by ten transformation configurations, plus once in a practice trial for each condition. We chose to use simple graphics rather than real photographs because we did not want subjects distracted by photographic details. We tried a variety of simple graphic images in our pilot testing and chose the four images displayed in Figure 4. Images 1, 2 and 3 are all straightforward to align. Image 4 is more difficult: the images can be aligned in a variety of configurations such that the rotated oval appears 'complete'. Other graphics tested in our pilot phase exhibited similar properties. Although we were interested in examining how subjects fared under the different conditions for more difficult alignment tasks, we only included one 'difficult to align' source image. We wanted to keep the complete experiment session time under one hour and minimize subject frustration.

T	S	R	Image	T	S	R	Image
0.0	1.2	0	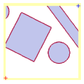	0.0	1.1	0	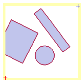
-0.1	0.9	0	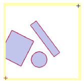	-0.2	0.8	0	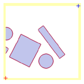
-0.15	1.0	10	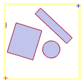	-0.1	1.0	5	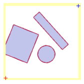
0.0	1.0	-5	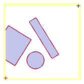	0.0	1.0	-10	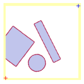
0.0	1.2	-10	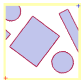	-0.1	0.9	-5	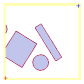

Table 1: Transformation configurations used in the experiment and displayed for one of the partial images. T is x-axis translation factor, S is scale factor and R is rotation factor in degrees.

Transformation Configurations In a typical virtual image alignment situation, the images are similar in alignment. If one of the images is rotated 90, 180 or 270 degrees with respect to the other, there are usually buttons or menu operations that can automatically perform these large rotations. For this reason, we used only slight rotations (10 degrees or less). Similarly, images to be aligned are usually of similar, though not identical scale. Finally, there is usually some overlap between the images, requiring one image to be translated over the other image slightly. These small 'jitter' factors are meant to reflect the slight imperfections one might see in a set of digital images taken by an amateur photographer trying to capture a panoramic scene. For simplicity, we did not include non-linear perspective transformations – these will be investigated in future work. The ten transformation configurations are listed in Table 1, along with a partial image for each configuration.

Condition	Image Translation	Image Rotation and Scale
SING	Drag with main button pressed	Drag with right button pressed
ASYM	Drag left mouse with main button pressed	Drag right mouse with main button pressed
SYM	Unified manipulation: drag both mice with main buttons pressed	

Table 2: Interaction techniques tested in the experiment.

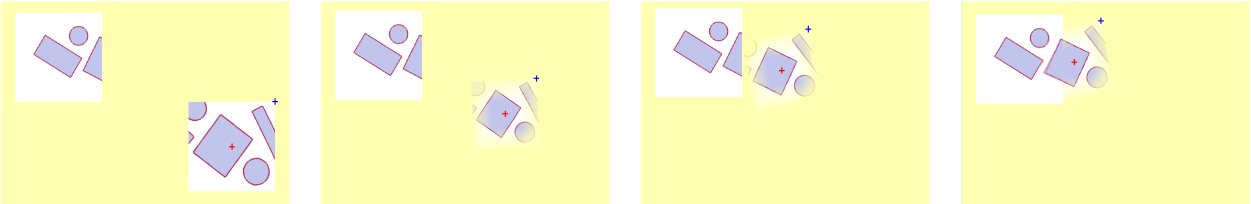

Figure 5: In the ASYM condition, the image can be translated by the non-dominant hand (red cursor) while being scaled and rotated by the dominant hand (blue cursor).

Conditions

Three conditions were tested in this experiment: single-mouse (SING), dual-mouse asymmetric (ASYM) and dual-mouse symmetric (SYM). Each condition represents a different image manipulation technique (see Table 2). For all three conditions, the cursor or cursors were locked to Image B.

SING Condition Our single mouse technique is similar to the image dragging technique in PTgui and Photoshop. Using the DH, the subject can press down on the main mouse button and drag the mouse to translate the image. Pressing the right mouse button instead allows the subject to drag the corner of the image (this differs slightly from the commercial software where the mode-switch to rotation is activated by clicking on a toolbar button). Dragging away from or towards the center of the image scales the image up or down, maintaining the aspect ratio (see Figure 6). Dragging in a circular motion around the center of the image rotates the image (see Figure 7).

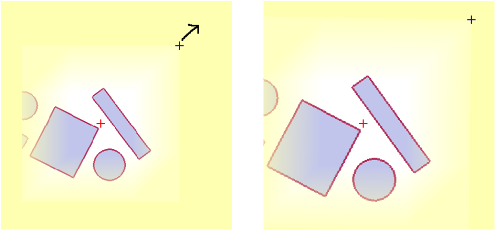

Figure 6: In the ASYM and SING conditions, the right cursor moves towards or away from the center of the image to scale it.

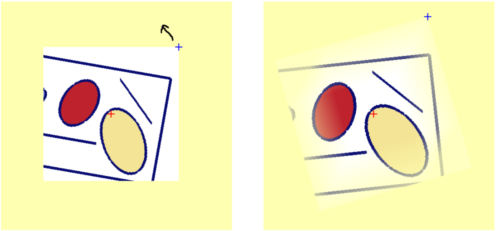

Figure 7: In the ASYM and SING conditions, the right cursor moves in an arc around the center of the image to rotate it.

ASYM Condition There are no descriptions of dual-mouse asymmetric techniques for image manipulation described in the HCI or image registration literature. Four degrees of freedom are required for the image alignment task: one for translation in x, one for translation in y, one for rotation around a fixed point and one for uniform scale. As each mouse offers two degrees of spatial input, we must split these four sub-tasks across the two mice. It is clearly inappropriate to split the two degrees of translation across the two mice. Therefore, the only sensible way to split the four sub-tasks across the two mice is to have one mouse control both x and y translation, and have the other mouse control rotation and scale. Thus, we chose an asymmetric separation of tasks where the NDH (using a red cursor attached to the center of Image B) controls the translation of the image while the DH (using a blue cursor attached to the upper-right corner of Image B) controls the rotation and scale of the image. We believe that this separation of tasks best conforms to Guiard's guidelines for asymmetric interaction [11]. In order for the rotation and scale to be performed by a single mouse, they both must occur with respect to the same point, which in this case is the center of the image. Because the NDH is controlling the position of the image, the scale and rotation of the image performed by the DH is done in the frame of reference set by the NDH. Note that the separation into sub-tasks is identical in the SING and ASYM conditions: translation is one sub-task and rotation/scale is the other sub-task.

For scale, the DH must move the cursor away from the center of the image to make the image larger or move the cursor closer to the center of the image to make the image smaller (see Figure 6). The aspect ratio of the image is maintained as the image is scaled. For rotation, the DH must drag the corner of the image in an arc around the center of the image, and the image rotates in the direction of the drag (see Figure 7). Finally, because the two mice can be moved simultaneously, the image can be translated by the NDH while being rotated and scaled by the DH (see Figure 5).

SYM Condition The symmetric interaction technique for image alignment is simple and follows the brief description of the two-handed 'stretchies' technique in Kurtenbach's T3 paper [15]. The cursors are attached to opposite corners of the image. By moving both cursors, the image can be rotated, scaled and translated simultaneously (see Figure 8). For this to work, the initial aspect ratio of the image is calculated. Once the cursors start moving, the diagonal line between them is used to determine the scale of the image, and the position of the two cursors determines both the location and orientation of the image.

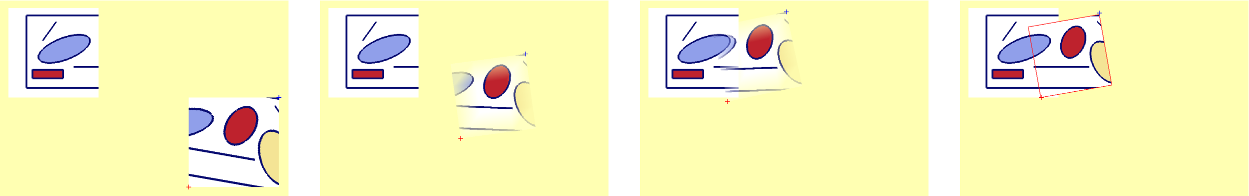

Figure 8: In the SYM condition, the image can be scaled, translated and rotated in a single fluid motion by moving both cursors around the screen simultaneously.

Performance Metrics

We measured trial completion time, reaction time and movement time. Additionally, we measured the time spent mode switching for the single-mouse technique. These measurements were calculated as follows:

- *Completion Time* was measured from the time the images were first presented until the images aligned correctly.
- *Response Time* was measured from the time the images were first presented until the first mouse movement.
- *Switch Time* for the SING condition was the accumulation of time periods measured from the time either mouse button was released until either mouse button was pressed.
- *Movement Time* was measured as completion time minus response time minus switch time (switch time only applies to SING condition).

In addition to performance times for each trial, we also measured the amount of parallelism exhibited by users in the SYM and ASYM techniques. We considered the subject to be using the two hands in parallel when both mice were moving continuously with the inner mouse buttons pressed.

Hypotheses

We have three hypotheses:

- H1: The mean completion times for SYM will be shorter than the mean completion time for ASYM which will be shorter than the mean completion time for SING.
- H2: The mean SYM completion time will be shorter than the mean SING completion time, even after removal of mode switch times from the SING trials.
- H3: The level of parallelism for the SYM condition will be higher than for the ASYM condition.

RESULTS

An analysis of the timing results for all trials across all subjects shows evidence to statistically support all three hypotheses. Practice trials were not included, so a total of (24 subjects × 3 conditions × 4 images × 10 transformation configurations =) 2880 trials were analyzed.

Mean Completion Times

Figure 9 shows the mean trial completion time under each condition as the height of each column. The mean completion time for SYM is 9.7 seconds, approximately 87% faster than for SING (18.3 seconds). The mean completion time for ASYM is 13.8 seconds or approximately 41% faster than the SING technique. This supports H1: subjects performed the task faster under the dual-mouse conditions than under the single-mouse condition, and faster under the SYM condition than under the ASYM condition. A repeated measures analysis of variance (RM ANOVA) for trial completion time gives an F-ratio for condition of 120.54 ($p_{2,1908} \leq 0.0001$), ensuring that the results in Figure 9 are statistically significant.

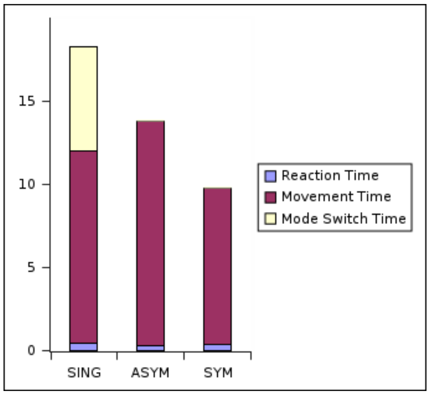

Figure 9: Mean trial completion times for each condition broken down vertically by reaction time, movement time and mode-switch time.

Cognitive Load Analysis

With the SING technique, time must be spent switching modes. This time was measured from the time one mouse button was released until another mouse button was pressed. When this mode-switching time is removed from the completion times for the SING condition, the movement time can be compared to the SYM and ASYM conditions. The mean movement times for the three conditions are shown as the purple sections of each column in Figure 9. Note also from Figure 9 that reaction time did not differ significantly across the three conditions. The mean movement times support H2: the SYM condition is significantly faster than the SING condition even after mode-switch times are removed. The RM ANOVA for movement time gives an F-ratio for condition of 59.50, ($p_{2,1908} \leq 0.0001$), showing that these results are statistically significant. The difference in mean movement times suggests that the shorter time seen in the symmetric technique is not simply because there is extra time taken in the single-mouse technique for mode switches. The mode switching takes an average of 6.3 seconds for each trial. However, there is still a 2.2 second difference in the movement times of the SYM condition and the SING condition. As well, there is a 4.1 second difference in the movement times of the SYM condition and the SING condition. These times suggest that the unified task representation in the SYM condition imposes a lower cognitive load than the other two techniques.

With mode-switching time removed, the SING condition is faster than the ASYM condition. Considering that the image registration task is split into the same sub-tasks in the SING and ASYM conditions (translation is one sub-task, and scale-rotation is the other sub-task), some other factor is producing statistically different average movement times (13.47 seconds versus 11.58 seconds) after the mode-switch time is removed. We postulate that the split into sub-tasks is a contributing factor for both conditions: SYM is 4.1 seconds faster than ASYM and 2.2 seconds faster than SING. Given these numbers, there is an approximate 2 second cost associated with the separation of the task into two sub-tasks. So, being able to rotate, scale and translate simultaneously saves 2 seconds per trial in the symmetric technique. However, there is still an extra 2 second cost unaccounted for in the ASYM technique compared to the SING technique. This may be an additional cognitive burden related to remembering which hand is doing which task. This isn't an issue with the SYM condition, because the two hands have exactly the same role.

Parallelism Analysis

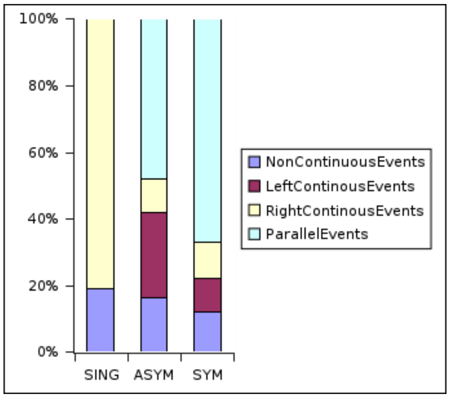

Figure 10: Percentage of events considered parallel, left mouse continuous, right mouse continuous and non-continuous for each condition.

We were interested in comparing the amount of parallel interaction between the SYM and ASYM events. When both hands *can* be used at the same time, the amount that they *are* used at the same time can be a good indicator of cognitive load. We measured parallelism by calculating the percentage of time that both mice were active simultaneously. To be considered active, a mouse event had to occur within two milliseconds of the previous event from the same mouse. For the two mice to be considered simultaneously active, the two mice both had to be active during the same time interval.

Table 3 shows the average percentage of mouse events per trial that occurred during parallel action under the two dual-mouse conditions. The breakdown of events into parallel, left-mouse continuous, right-mouse continuous and non-continuous is displayed in Figure 10. These results support H3: subjects engaged in more parallel activity under the SYM condition than under the ASYM condition. The RM ANOVA for percentage of parallel activity gives an F-ratio for condition of 1164.0 ($p_{1,954} \leq 0.0001$), showing that these results are significant. This also supports the conjecture that the ASYM condition imposes a higher cognitive load on the subject than the SYM condition.

Condition	% Parallelism	Std Dev
ASYM	48.2%	15.7%
SYM	67.1%	8.6%

Table 3: Percentage of mouse events considered parallel (occurring within 2 msec of previous events from both mice).

Learning Effects

Condition Order	mean total time	stdev
SYM, ASYM, SING	13.88	11.45
SYM, SING, ASYM	13.69	10.49
ASYM, SING, SYM	13.48	11.93
ASYM, SYM, SING	13.49	13.44
SING, SYM, ASYM	13.28	9.45
SING, ASYM, SYM	15.88	20.98

Table 4: Mean completion times (seconds) by condition order.

We separated the 24 subjects into six groups and counterbalanced the presentation of the three conditions across the groups. Table 4 shows the condition presentation order, as well as the mean total trial time for each group. There is little variance across the six groups, and this is confirmed by the repeated measures analysis for total trial time, which gives an F-ratio for presentation order of 2.12 ($p_{5,1908} = 0.06$), which is not statistically significant. The final group has the most obvious deviation, with a mean total time approximately 2 seconds longer than the other groups. On closer examination, the subject who had the highest mean total trial time was in this group. That subject's mean total trial time was 22.12 seconds, whereas the overall mean total time across all subjects was 13.95 seconds. Given the statistical insignificance of this differences across the groups, we conclude that the learning effects are minimal.

Image Set Effects

The four images that we chose for the experiment task are shown in Figure 4. Each of the four images was presented to each subject 30 times during the experiment. As discussed earlier, we expected image alignment on three of these images to be relatively easy, but we expected subjects to take longer to align Image 4. We anticipated that the completion times for Image 4 would exhibit larger differences between the conditions. As shown in Figure 11, the total completion times are longer trials using Image 4, and the difference is more pronounced for the SING technique, where the 21.27 second average completion time is approximately 3 seconds longer than the completion times for the other images. The difference for both the SYM and the ASYM techniques is approximately 1 second. This suggests that the performance

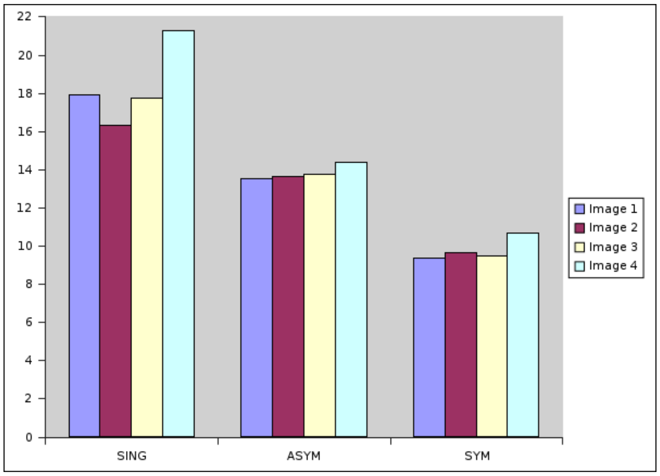

Figure 11: Mean completion times (seconds) by condition and image.

benefits from using dual-mouse techniques in this task become more pronounced with task difficulty.

Subjective Evaluation

After completing the experiment, subjects were asked to rank the three interaction techniques. Specifically, they were asked, “Imagine you had a job lining up digital images all day. Rank the three techniques in order of preference according to which you would like to use in your job.” The subjects were not told of their performance with any of the techniques before answering this question. The symmetric technique was preferred by 22 of 24 candidates. One of the 24 subjects preferred the single-handed technique and one of the subjects preferred the asymmetric technique (although both performed fastest with the symmetric technique). This preference ranking further supports our assertion that the task is naturally symmetric.

CONCLUSIONS AND FUTURE WORK

Previous work in two-handed interaction techniques has often focused on asymmetric interaction, following the guidelines set out by Guiard. We have demonstrated that a symmetric technique for simultaneous rotation, translation and scale yields significant performance benefits over both a dual-mouse asymmetric technique and the standard single mouse technique in an image alignment task. The benefits of the symmetric technique over the single-mouse technique are statistically significant even after the mode-switch time is removed, suggesting that the performance benefits are not only due to motor timing, but to cognitive benefits from a unified task representation. Additionally, the level of parallelism demonstrated by subjects was higher for the symmetric technique than for the asymmetric technique, which also points to the benefits of a unified task representation. This work contributes a formal validation of a symmetric interaction technique, which is important because evaluation of symmetric interaction is lacking in the HCI literature. User interface designers should not automatically assume that all two-handed interaction should be asymmetric.

One aspect of image alignment that was not part of this evaluation is the skewing of images to reflect lens angle changes (non-linear perspective transformations). We plan to implement and test a symmetric technique in which the two mice can choose same-side corners of an image and skew the image by pulling the cursors apart or pushing them together, while leaving the opposite edge of the image anchored. Additionally, we want to see how the results reported here generalize to other tasks involving rotation, translation and scale. We plan to investigate symmetric interaction in other domains, such as in windowing systems.

REFERENCES

1. Ravin Balakrishnan and Ken Hinckley. Symmetric bimanual interaction. In *CHI 2000 Proceedings*, 2000.
2. Michel Beaudouin-Lafon. Novel interaction techniques for overlapping windows. In *UIST 2001 Proceedings*, pages 153–154. ACM, 2001.
3. W. Buxton and B. Myers. A study in two-handed input. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 321–326. ACM Press, 1986.
4. House Internet Services B.V. PTgui. `http://www.ptgui.com/`.
5. Didier Casalta, Yves Guiard, and Michel Beaudouin-Lafon. Evaluating two-handed input techniques: Rectangle editing and navigation. *ACM CHI'99 Conference (Extended Abstracts)*, pages 236–237, 1999.
6. Stephane Chatty. Extending a graphical toolkit for two-handed interaction. In *UIST'94 Proceedings*, pages 195–204, 1994.
7. Stephane Chatty. Issues and experience in designing two-handed interaction. In *CHI'94 Proceedings*, 1994.
8. Chia-Yen Chen and Reinhard Klette. Image stitching - comparisons and new techniques. In *Lecture Notes in Computer Science*, volume 1689, pages 615–622. Springer-Verlag GmbH, Jan 1999.
9. PanaVue Company. Panavue image assembler. `http://www.panavue.com/`.
10. George Fitzmaurice, Hiroshi Ishii, and William Buxton. Bricks: Laying the foundations for graspable user interfaces. In *CHI'95 Proceedings*, 1995.
11. Yves Guiard. Asymmetric division of labour in human skilled bimanual action: The kinematic chain as a model. *Journal of Motor Behaviour*, pages 486–517, 1987.
12. Ken Hinckley, Mary Czerwinski, and Mike Sinclair. Interaction and modeling techniques for desktop two-handed input. In *UIST'98 Proceedings*, 1998.
13. Paul Kabbash, William Buxton, and Abigail Sellen. Two-handed input in a compound task. In *CHI'94 Proceedings*, 1994.

14. Russell Kruger, Sheelagh Carpendale, Stacey D. Scott, and Anthony Tang. Fluid integration of rotation and translation. In *CHI'05 Proceedings*, 2005.

15. Gordon Kurtenbach, George Fitzmaurice, Thomas Baudel, and Bill Buxton. The design of a GUI paradigm based on tablets, two-hands and transparency. In *CHI'97 Proceedings*, 1997.

16. Celine Latulipe. Symmetric interaction in the user interface. In *Companion Proceedings of UIST 2004*, pages 63–66, 2004.

17. Andrea Leganchuk, Shumin Zhai, and William Buxton. Manual and cognitive benefits of two-handed input. *ACM Transactions on Computer-Human Interaction*, 5(4):326–359, 1998.

18. Ignacio Llamas, Byungmoon Kim, Joshua Gargus, Jarek Rossignac, and Chris D. Shaw. Twister: a space-warp operator for the two-handed editing of 3D shapes. *ACM Trans. Graph.*, 22(3):663–668, 2003.

19. I. Scott MacKenzie and Yves Guiard. The two-handed desktop interface: are we there yet? In *CHI '01 extended abstracts on Human factors in computer systems*, pages 351–352. ACM Press, 2001.

20. Russell Owen, Gordon Kurtenbach, George Fitzmaurice, Thomas Baudel, and Bill Buxton. When it gets more difficult, use both hands - exploring bimanual curve manipulation. In *GI 2005 Proceedings*, 2005.

21. Roope Raisamo and Kari-Jouko Raiha. Techniques for aligning objects in drawing programs. In *UIST'96 Proceedings*, 1996.

22. Chris Shaw and Mark Green. THRED: A two-handed design system. *Multimedia Systems*, 5:126–139, 1997.

23. Rick Szeliski. Image alignment and stitching. Technical Report MSR-TR-2004-92, Microsoft Research, 2004. `http://www.csie.ntu.edu.tw/~cyy/courses/vfx/papers/Szeliski2005IAS.pdf`.