Interacting with Stroke-Based Rendering on a Wall Display

Jens Grubert,[1] Mark Hancock,[1] Sheelagh Carpendale,[1] Edward Tse,[1] Tobias Isenberg[1,2]

[1]University of Calgary
Canada
{jgrubert|msh|sheelagh|tsee}@cpsc.ucalgary.ca

[2]University of Groningen
The Netherlands
isenberg@cs.rug.nl

ABSTRACT
We introduce two new interaction techniques for creating and interacting with non-photorealistic images using stroke-based rendering. We provide bimanual control of a large interactive canvas through both *remote pointing* and *direct touch*. Remote pointing allows people to sit and interact at a distance with an overview of the entire display, while direct-touch interaction provides more precise control. We performed a user study to compare these two techniques in both a controlled setting with constrained tasks and an exploratory setting where participants created their own painting. We found that, although the direct-touch interaction outperformed remote pointing, participants had mixed preferences and did not consistently choose one or the other to create their own painting. Some participants also chose to switch between techniques to achieve different levels of precision and control for different tasks.

Author Keywords
Wall display, Nintendo Wii Remote (Wiimote) and Nunchuck, direct touch (DT), non-photorealistic rendering (NPR), stroke-based rendering (SBR).

ACM Classification Keywords
H.5.2 [Information Interfaces and Presentation]: User Interfaces—Interaction styles, Input devices and strategies; I.3.m [Computer Graphics]: Miscellaneous—Non-photorealistic rendering.

(a) Remote pointing with Wiimote.

(b) Direct touch using SMART's DViT.

Figure 1: Remote pointing and direct-touch interaction with the Interactive Canvas.

INTRODUCTION
The physical world provides artists with the freedom to achieve various levels of precision and control. On the small scale, they use a variety of brushes, palette knives, and surrounding objects to apply paint according to their current inspiration. On the large scale, they not only use small wrist motions, but often use full-body movements to achieve different effects. Many automatic tools now exist to alter digital images. These tools are typically designed to be used inside a computer program and so are manipulated with standard input devices, such as mice or styli. While many creative interaction techniques have been designed with these devices, the limited freedom of movement can make it difficult to achieve the same expressive power as in the physical world.

To begin to address the need for incorporating these types of freedoms into digital painting, we introduce new bimanual interaction techniques (Figure 1). Our techniques combine Nintendo's Wii controller with a large direct-touch canvas and provide support for the creation and manipulation of non-photorealistic paintings. These techniques allow people to sit or stand as they desire, and to interact directly with the canvas or to point to it from a distance.

We performed a user study to compare direct-touch interaction to remote pointing in both a controlled setting with timed tasks and an exploratory setting in which participants created a digital painting. Our results suggest that, while direct touch may be more efficient, many other issues make remote pointing a useful alternative. We also observed that participants tended to sacrifice accuracy in remote pointing to achieve speeds close to those of direct touch. By trading off accuracy for speed, people could reap the benefits of remote pointing, such as an overview of the entire display and the ability to sit and rest while performing small motions instead of large reaching movements.

We first review related work and then introduce the interaction methods that we created in the context of the Interactive Canvas application. Then, we present and discuss the results of a user study that compares our two bimanual interaction techniques.

RELATED WORK

We review literature in five related areas: direct vs. indirect interaction, bimanual interaction, freehand pointing on large displays,interaction with non-photorealistic rendering, and interaction with digital painting.

Direct vs. Indirect Interaction. There are a number of studies comparing the use of direct-touch interaction to indirect mouse input. These studies focus primarily on the performance (i. e., speed and accuracy) of each input technique. The seminal work of Card et al. [4] found that mouse input performance compared quite favourably to direct stylus input. Sears and Shneiderman [21] compared the performance of mouse input to touch screen input in a unimanual task. They found that for targets 6.4 mm or more in width, touch screen selection was faster than mouse selection. Further, for targets 12.8 mm in width, touch screen selection resulted in about 66% fewer errors. However, even with the apparent superior performance of direct-touch input, participants still preferred mouse input. The authors attributed this disparity to arm fatigue when using an upright direct touch display. Since people have to raise their arms to interact with the touch screen, arm fatigue arises when working over long periods of time. Ahlström et al. [1] demonstrated that changing the screen mounting angle (e. g., a 45° drafting table orientation) can substantially reduce fatigue over large displays.

Forlines et al. [5] compared the performance of using mice and direct touch on a large digital table for both one- and two-handed interaction. Their results indicate that the users may be better off using a mouse for unimanual input and their fingers for bimanual input when working on a large horizontal display. While direct touch did not lead to greater performance in terms of speed and accuracy for unimanual tasks, the authors suggested that direct touch may still be beneficial for other characteristics such as spatial memory and awareness of others' actions in a multi-user setting.

Closely related to our work, Myers et al. [16] compared the use of laser pointers to both mouse and direct-touch interaction with a SmartBoard. Direct touch was found to be the fastest and most accurate technique, followed by the mouse and, finally, laser pointers. Similar to previous studies, user responses indicated that the mouse was preferred over direct touch. The authors argued that jitter in the laser input device affected the accuracy of the input and, despite this limitation, laser pointers were still beneficial for the convenience of not having to walk up to touch the display.

Bimanual Interaction. There has been significant research in the area of bimanual interaction, in terms of theory [7], empirical studies [3], and interaction design [3, 9]. Some of the earliest work in the HCI field on bimanual interaction is the study by Buxton and Myers [3], which clearly articulated the benefits of bimanual input on graphical user interface tasks. They showed benefits for leveraging the non-dominant hand for reference changes such as scrolling while using the dominant hand for precise selection. Guiard [7] described a theoretical model for understanding the nature of this bimanual action called the Kinematic Chain. The non-dominant hand remains near the root of the kinematic chain and can be used for coarse hand movements while precise selection is achieved lower on the kinematic chain through the dominant hand. Hinckley et al. [9] found that performance was significantly reduced when these roles were reversed. This suggests that the dominant hand operates relative to the frame-of-reference of the non-dominant hand.

Freehand Pointing on Large Displays. Ray casting is a commonly used technique for pointing to distant objects on a large display (e. g., [16, 17, 26]), where the cursor is drawn as the intersection of the ray from the hand/pointer and the screen. Laser pointers are obvious candidates for implementing ray casting, and many people have explored how they can be used. Myers et al. [16] considered different laser pointer form factors (pen, glove-mount, scanner, toy gun) to see how they minimized hand jitter and affected aiming. Parker et al.'s *TractorBeam* [17] affords selection on a tabletop display by having people point the tip of the six degree-of-freedom pen at distant targets. Other ray casting devices include data gloves, Nintendo's Wiimote, wands tracked by motion capture systems, etc.

Interaction with Non-Photorealistic Rendering. Much of the work in non-photorealistic rendering (NPR) has focused on the automatic creation of imagery [6, 25]. The subset of NPR most closely related to our work is stroke-based rendering (SBR, [8]) and specifically the painterly rendering technique Interactive Canvas [20]. This technique and many other approaches strive for interactive rendering. For most NPR techniques this means that images are generated at interactive or real-time frame-rates using a predefined set of parameters, as opposed to interaction during the image creation. A few exceptions explore the possibilities for user interaction during the rendering. For example, WYSIWYG-NPR [12] allows users to interactively stylize automatically extracted strokes; the Interactive Canvas [20] focuses on interactive placement of strokes and manipulation of their properties based on spatial interaction buffers [10]. In our work, we examine possible interaction techniques that make use of the Interactive Canvas paradigm and compare two different interaction approaches in a user study.

Interaction with Digital Painting. Digital painting is a major form of artistic depiction today. In pixel-based approaches (e. g., Adobe Photoshop or Gimp) users typically interact using a paintbrush metaphor. In contrast, vector-

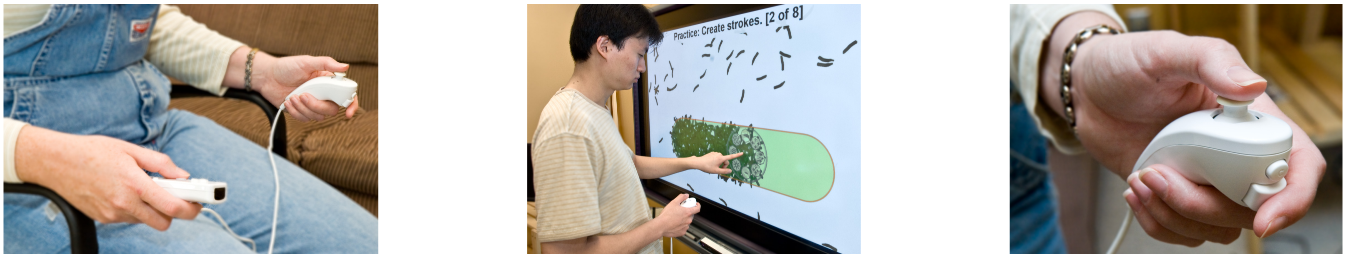

(a) Remote pointing (b) Direct-touch (c) Nintendo Nunchuck

Figure 2: Two new techniques for painting on the Interactive Canvas include (a) remote pointing and (b) direct touch interaction. Both techniques use the Nintendo Nunchuk (c).

graphics techniques (e. g., Corel Draw and Adobe Illustrator) concentrate on changing attributes of primitives rather than manipulating pixels on a grid. The Interactive Canvas paradigm can be seen as benefiting from these two extremes as it uses a paintbrush metaphor to change the properties of primitives rather than to manipulate pixels. Interaction with the traditional digital painting programs is typically provided with a mouse, but often people employ touch-sensitive tablets to gain better control due to added pressure sensitivity and a more brush-like interaction. Researchers interested in interacting with non-photorealistic rendering have also considered the painting metaphor. They have developed painting systems that simulate physical brushes [2], brushes that capture and use real-world textures [18], or real brushes that are tracked in physical space to manipulate a digital painting [14]. Other techniques employ three-dimensional painting [13, 19], use capacitive tracking on tables to interact with simulated fluid jet painting [15], or even use facial expression recognition to control parameters of an automatic digital painting based on the emotional state of a viewer [24].

BIMANUAL CONTROL FOR THE INTERACTIVE CANVAS

In contrast to automatic and parameter-tweaking approaches for the creation of non-photorealistic paintings, we extend the basic approach of the Interactive Canvas by developing a bimanual interface.

Interactive Canvas

The Interactive Canvas [20] is in itself unique as a digital painting approach, because it combines visual richness of pixel-based painting with some of the freedoms of vector-based painting. Similar to pixel-based approaches, images are created with fully-rendered primitives, providing such attributes as texture and shading in each brush stroke. However, since the Interactive Canvas offers modeling with rendering primitives, all primitives remain interactive throughout the creative process. A rendering primitive can be defined as any small piece of an image such as different types of brush strokes (lines, points, dabs) and any other image or image component such as images of teapots, popcorn, or leaves. These primitives can be used as basic elements from which a painting is constructed and adjusted. Throughout the creative process one can continuously create, take apart, reassemble and adjust these rendering primitives interactively through local and direct manipulations, leveraging the system's immediate visual feedback. In addition, the primitives have a certain size larger than a single pixel and thus introduce abstraction, resulting in expressiveness of the created images. While this interactive NPR technique holds potential for engaging people while creating artwork, the new freedoms open up many questions about the type and style of interactions that should be developed.

Interface Design

To make use of the functionality offered by the Interactive Canvas and to leverage its potential, we examined several methods of interacting with the virtual canvas. Specifically, we explored the use of both table and wall displays, direct-touch interaction, speech commands, bimanual controls, and indirect pointing. Although we could provide the functionality with many different combinations of these interface components, two bimanual combinations stood out as particularly promising interaction techniques: remote pointing and direct-touch interaction (see Figure 1).

While both remote pointing and direct touch have a natural mapping to the available functionality, it was not clear whether people would prefer to develop their digital paintings while working directly with touch on the screen or while comfortably sitting a few feet away and interacting via remote pointing. To answer this question we developed the basic required functionality for both setups and integrated them so that a person could modelessly switch from distance pointing to direct touch at will.

Interaction Techniques

Both methods we provide to create and interact with paintings—remote pointing and direct touch—use bimanual control (Figure 2). The non-dominant hand is used to indicate what effect the dominant hand interaction will have. In both techniques, the non-dominant hand is used to control a Nintendo Nunchuk (built for the Wii gaming console), shown in Figure 2(c). This device provides a small joystick that can be controlled comfortably with the thumb. This joystick is positioned inside an octagon, and thus has eight discrete and easily acquired positions. The Nunchuk also has two buttons on the front (buttons *C* and *Z*). We use the eight joystick positions together with the centre resting position for nine primary functions of the system (shown in Figure 3). We allow control of the size of the area affected by the dominant hand by holding button *Z* and moving the joystick up or down. We provide control of the rate of change by the dominant hand through a similar interaction with button *C*. Holding button

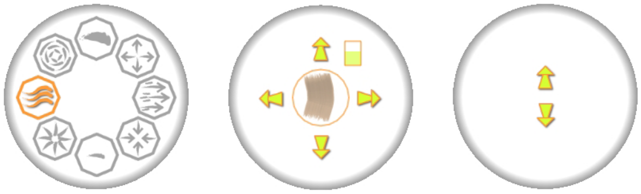

(a) Tools menu (align selected) (b) Strokes menu (c) Radius control

Figure 3: On-screen menus for selecting actions.

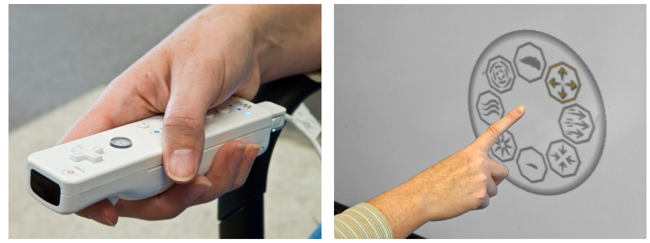

(a) Wiimote. (b) Direct-touch.

Figure 4: Remote pointing using the Wiimote vs. direct-touch interaction directly on the wall display.

C and moving the joystick left or right toggles through a selection of alternate strokes that can be used. Strokes can be removed by holding both buttons *C* and *Z* while interacting with the dominant hand. In all cases, the user holds the joystick in the appropriate direction throughout the dominant-hand interaction. We chose this requirement because it has been shown before that forcing the user to *maintain* this position will prevent mode errors [22].

Interaction at a Distance

In order to invoke dominant-hand interactions at a distance, the user can point to the display with the Nintendo Wii Remote controller (Wiimote; see Figure 4(a)) and press a trigger (button *B*). The Wiimote is a wireless device that, among other data elements, provides x and y coordinates on the display. The coordinates are derived by using a camera in front of the controller that is tracking the positions of stationary infrared LEDs that are placed below or above the display (for more details see, e. g., [23]). We chose to invoke actions with the trigger button *B*, because we assumed that it would cause less drift of the cursor than other buttons and it was also specifically designed for triggering actions.

Direct-Touch Interaction

Users can also invoke dominant-hand interactions by directly touching the display (Figure 4(b)). Since the Wiimote is no longer needed for direct touch, it can be placed in a pocket or on a belt, while still allowing freedom of movement with the Nunchuk (Figure 1(b)).

USER STUDY

We performed a user study to evaluate the interaction techniques that we developed for creating NPR paintings. In this evaluation, we were interested both in the performance of the techniques when doing simple tasks, and how people would use them to create an entire painting.

Participants

We recruited sixteen paid participants from the area of a local university (seven male, nine female). Ages ranged from 21 to 66 years ($M = 30.3$ years, $Mdn = 28$ years, $SD = 10.2$ years). Five had used a Wiimote before, all were right-handed (with one claiming to have converted from ambidextrous at birth), and seven had an artistic background.

Apparatus

Participants performed the experiment at a plasma wall display with a resolution of 1360×768 pixels and a display area of $135\,\text{cm} \times 76\,\text{cm}$, mounted so that its bottom was 106 cm off of the ground. Direct-touch input was provided through SmartBoard DViT technology and remote pointing through a Nintendo Wiimote and Nunchuk. For the direct-touch condition, participants were asked to stand directly in front of the display with the Nunchuk in one hand. For the remote-pointing condition, participants were asked to sit in a chair that was 46 cm high and placed 165 cm in front of the display (eye-to-display distance approx. 195 cm), with the Nunchuk in one hand and the Wiimote in the other. The infrared LED markers used to detect the Wiimote's position were placed at the bottom of the screen.

Procedure & Design

The user study consisted of two phases. In the first phase, we were interested in measuring the performance of each technique as the participants performed controlled painting tasks. In the second phase, we were interested in observing the participants' behaviour when they were given the freedom to choose how to interact.

Phase I: Controlled Tasks

In this phase of the experiment, we had participants perform the following four tasks:

- create strokes (create),
- align strokes horizontally (align),
- orient strokes in a star shape (star), and
- repel strokes in a circular pattern (repel).

These four tasks can be invoked with the Nunchuk joystick (Figure 3(a)) using the centre rest position (create), the left position (align), the bottom-left position (star), and the top-right position (repel). While holding the correct position, the participant then touched the display (in the direct-touch condition) or pointed with the Wiimote and pressed the *B* button (in the remote-pointing condition) to invoke the action.

Each participant was asked to perform five blocks of trials for each of the two techniques. Each block consisted of 20 trials (5 repetitions of each task) for a total of 200 trials (2 techniques $\times$ 5 blocks $\times$ 20 trials). For each technique, participants began with a practise block of 8 trials (2 trials per task) and were reminded that they could take a break after each block. For each trial, the instruction (i. e., which task to perform) was displayed at the top of the screen and a target area was displayed in the centre. The participant was asked to perform the described task inside the target area

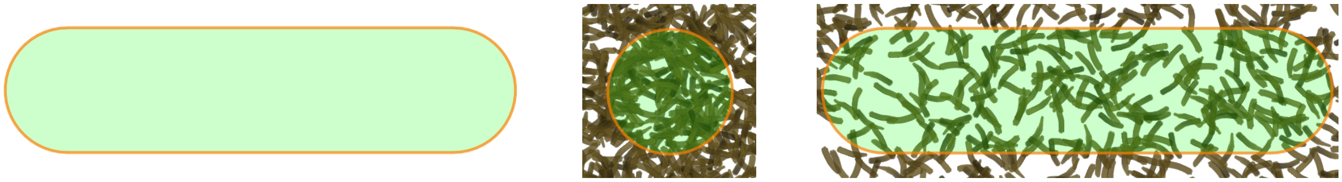

(a) For create task. (b) For star and repel tasks. (c) For align task.

Figure 5: Target areas for the tasks in Phase I of of our study.

as quickly as they could, but to affect the area outside the boundary as little as possible.

For the create and align tasks, the target area was a long horizontal oval (Figures 5(a) and 5(c)) and for the star and repel tasks, the target area was a circle (Figure 5(b)). For each trial, distractor strokes were distributed randomly outside the target area. For the align, star and repel tasks, a more dense concentration of strokes was distributed in an area with double the height of the target area (centred at the same location), providing a dense set of strokes on which to perform the task. The participant indicated that they were finished each trial by pressing the *Z* button on the Nunchuk, which also started the next trial in the block.

The area affected by a touch in the direct-touch condition and by the Wiimote in the remote pointing condition was a circle the same height as the target area. Thus, the ideal movement was to draw a straight line in the create and align tasks and to acquire and dwell on a target at the centre of the circle in the star and repel tasks.

Phase II: Painting

In addition to the tasks from Phase I of the experiment, the participant was introduced to the following additional functionality. They could also:

- orient strokes radially,
- move strokes inward—like a black hole,
- move strokes along a line,
- make strokes larger,
- make strokes smaller,
- adjust the size of the affected area,
- alter the stroke type,
- adjust the rate at which strokes were created or erased, and
- erase strokes.

Participants were shown four examples of paintings created with our system along with the template images used to create them. They were then asked to create their own painting based on one of four photographs (Figure 6) using the provided tools. This photograph was used to automatically determine the colour of each stroke based on its position in the canvas; participants were thus only required to alter the number, size, orientation, position, and type of the strokes.

Hypotheses & Focus

The first phase of the experiment was designed as a hypothesis test to compare direct touch to remote pointing, specifically in the context of the tasks. We were also interested in

Figure 6: Template images, one of which each participant was asked to interpret in the painting phase.

how participants learned to use the devices over time. We thus had the usual null hypotheses associated with our factorial design.

The second phase of the experiment was designed to provide the opportunity to observe our system in use. Our focus was on the following aspects of this interaction:

- the choice of interaction technique,
- whether participants would switch between interaction techniques,
- what tools the participants would choose to use,
- whether participants would rate certain aspects of the system particularly enjoyable or frustrating, and
- whether participants would enjoy working with the system in general.

Data Collection

Participants were videotaped and the experiment was followed by an informal interview in which they were asked to comment on ease of use, problems encountered, and overall opinions for each of the techniques. Timing data and input device coordinates (for both direct-touch and the Wiimote) as well as the final size, orientation, and position of each stroke were logged.

RESULTS & DISCUSSION

In the first phase of the experiment, we were interested primarily in performance, so that we could focus on observing behaviour in the second phase. We thus present results for speed and accuracy for the first phase only. Data were analysed using a within-participants analysis of variance for the following three factors:

- block (1–5),

- task (create, align, star, repel), and
- device (direct touch, remote pointing).

We adjusted significance values for all post-hoc pairwise comparisons using the Bonferroni correction.

Speed

We analysed the task completion times (TCT) for each trial. TCT includes several components including: the time to read the instruction, the time to react, the time to select the appropriate action with the Nunchuk, the time to acquire the target, and the time to perform the action. We also separately analysed the time to perform the action, but we report only our results for TCTs, as the effects, interactions and mean differences were similar.

Factor	F-score	p-value
device	$F(1,13) = 8.7$	$p = .01$
block	$F(4,52) = 46.3$	$p < .001$
task	$F(3,39) = 5.8$	$p < .01$
device × block	$F(4,52) = 0.6$	$p = .66$
device × task	$F(3,39) = 5.2$	$p < .01$
block × task	$F(12,156) = 1.4$	$p = .16$
device × block × task	$F(12,156) = 2.4$	$p < .01$

Table 1: ANOVA results for task completion times.

We summarize the main effects and interactions in Table 1. The main effect of *device* shows that participants were significantly faster with direct touch ($M = 4.20$ s, $SE = 0.33$ s) than with remote pointing ($M = 5.32$ s, $SE = 0.49$ s). The main effect of *block* reflects an expected learning effect; pairwise comparisons showed that participants were significantly slower in block one than all future blocks ($p < .05$), and significantly slower in block two than blocks three and five ($p < .05$), but that differences between blocks three to five were not significant ($p > .05$). For the main effect of *task*, post-hoc tests showed that participants were significantly slower in the align task than in the star ($p < .01$) and repel ($p < .01$) tasks, but no other pair of tasks were significantly different ($p > .05$). We suspect that the align task was slower because, although we observed quick movement for this task, participants sometimes needed to "correct" the result with a second or third pass.

The interaction between *device and task* (see Figure 7) shows that the difference in performance for the align task

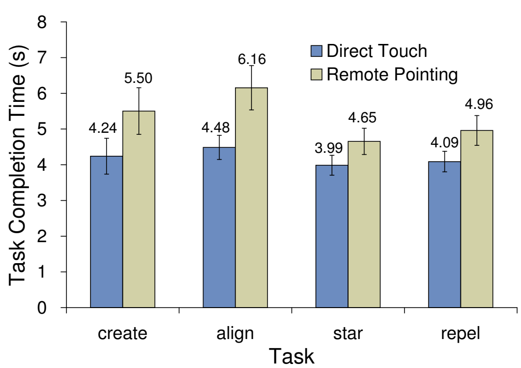

Figure 7: Interaction between device and task.

Task	Block 1	2	3	4	5
create	$p = .19$	**p < .01**	**p < .01**	**p < .01**	**p < .001**
align	$p = .06$	$p = .05$	**p < .01**	**p = .03**	**p < .01**
star	$p = .20$	**p = .02**	$p = .17$	$p = .09$	$p = .62$
repel	**p < .01**	**p = .02**	$p = .49$	**p < .01**	$p = .35$

Table 2: Pairwise significant differences between devices.

depends on which device the participant used. That is, for remote pointing, the align task was significantly slower than the star ($p < .01$) and repel ($p = .01$) tasks and no other pairs were different (similar to the main effect of task), but no task pairs were different for the direct-touch condition ($p > .05$). The *three-way interaction* further illustrates these differences (see Figure 8). In addition, Table 2 shows the pairwise significant differences between devices for each task and each block. All mean differences show that direct touch was faster than remote pointing. These differences suggest that, for tasks requiring movement along a line (create and align), the improvement over time was greater for direct touch; but for tasks requiring only pointing and dwelling (star and repel), the improvement over time was greater for remote pointing. Note also in the latter case that the remote pointing improved to be not significantly different than direct touch by the final block.

Accuracy

We analysed two measures of accuracy: average distance (D_{avg}) and coordination (C). The average distance is a measure of how closely the participant matched the optimal trajectory. The coordination measure is the same as that used by Zhai and Milgram to measure coordination in six degree of freedom input devices [27], but cannot be calculated for the *star* and *repel* tasks, since the optimal path length is zero. For both of these measures, a lower value indicates higher accuracy with either the participant's finger or the Wii Remote. We define these two measures as follows:

$$D_{avg} = \frac{1}{|P|} \sum_{p \in P} distance(p, L_{opt})$$

$$C = \frac{length(P) - length(L_{opt})}{length(L_{opt})}$$

Where P is the set of points defining the path traversed during a trial (the touched points in the direct touch condition and the points traversed while holding the B button in the indirect condition) and L_{opt} is the optimal path for a trial (a line in the create and align tasks, a single point in the star and repel tasks).

Distance. There was a significant main effect of *device* ($F(1,14) = 172.3$, $p < .001$). Participants activated points closer to the optimal path with the direct-touch technique ($M = 7.8$ pixels, $SE = 0.4$ pixels) than with the remote pointing technique ($M = 14.1$ pixels, $SE = 0.6$ pixels). There was also a significant interaction between *device and task* ($F(3,42) = 3.9$, $p = .01$). Post-hoc comparisons showed that, for the direct-touch condition, the average distance to the optimal line was significantly closer for the create task than for the align task ($p = .01$), but no other pair of tasks was significantly different for either technique ($p > .05$).

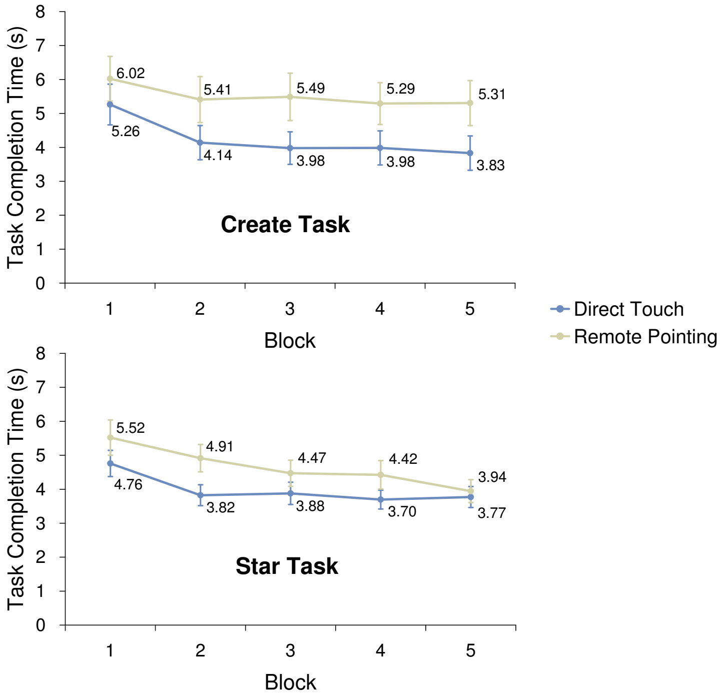

Figure 8: Three-way interaction between device, block, and task.

This isolated difference may be due to the fact that the create task invited more precision than the align task (which both required movement along a line), which was only achievable using the direct-touch device (see Figure 9). There were no other main effects or interactions ($p > .05$).

Coordination. There was a significant main effect of *device* ($F(1,14) = 8.6,\ p = .01$). Participants were more coordinated when using direct touch ($M = 0.25$, $SE = 0.02$) than when using remote pointing ($M = 0.51$, $SE = 0.11$). There was also a significant main effect of *task* ($F(1,14) = 14.8$, $p < .01$) as participants were more coordinated in the create task ($M = 0.24$, $SE = 0.03$) than in the align task ($M = 0.53$, $SE = 0.10$). There was also a significant interaction between *device and task* ($F(1,14) = 8.6$, $p = .01$). Post-hoc analysis showed that, in the align task, participants were significantly more coordinated with direct touch than with remote pointing ($p = .01$), but in the create task, this difference was not significant ($p = .16$). There were no other significant main effects or interactions ($p > .05$).

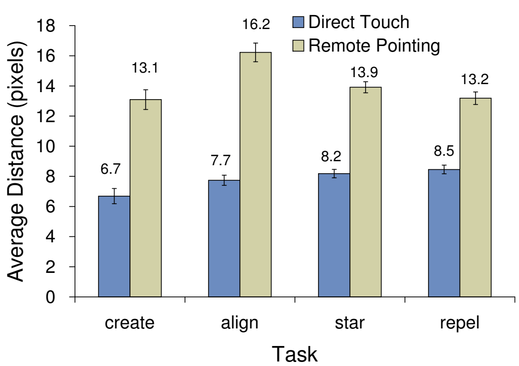

Figure 9: Average distances by task for both devices.

With the lack of significant differences involving the block factor, it appears that coordination is not affected by learning or fatigue (within the hour-long time frame of our study). Our results also suggest that coordination difficulties with remote pointing depend on the task. That is, in the align task, participants were less coordinated with remote pointing, but not so in the create task.

Questionnaire Data

We used a seven-point Likert scale for questions about speed, control, and expectation (Figure 10) and asked participants to state which device (if any) they preferred (Figure 11) for each task. They agreed that both devices were fast and responded as expected. This agreement was slightly stronger in the create and align tasks and overall for direct touch. Participants disagreed with direct touch being difficult to control in all tasks. For remote pointing, they disagreed with this statement for the star and repel tasks, but agreed for the align task and were neutral for the create task and overall. They showed a clear preference for direct touch, particularly for the create and align tasks. Note that participants were asked specifically about speed and control, but often commented that they preferred remote pointing for other reasons.

OVERALL DISCUSSION

In this section, we elaborate on our findings in Phase I and discuss them in terms of the observations we made in both phases. Consistent with previous findings, our results suggest that direct touch is faster and more accurate than remote pointing. These results alone suggest that direct touch is a better design choice than remote pointing; however, our observations point to a variety of other factors that may make remote pointing or a combination of both a better choice in practise. The importance of these other factors is reflected

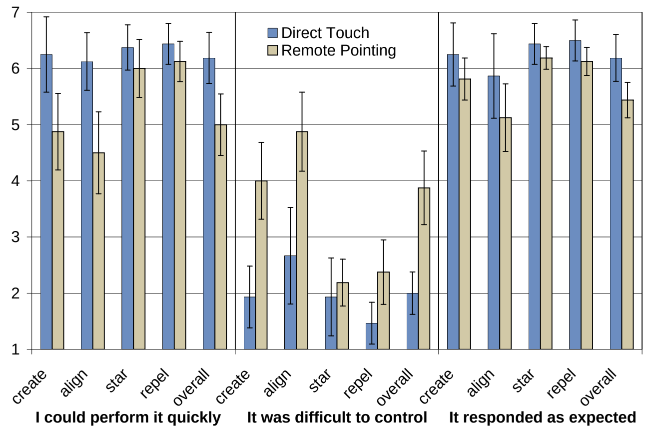

Figure 10: Participant responses for speed, control, and expectation (1 = strongly disagree, 7 = strongly agree).

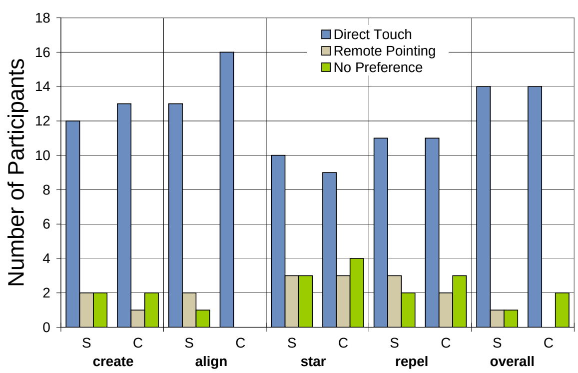

Figure 11: Participant preferences for speed and control.

by the fact that, in the second phase, only seven participants chose to use direct touch, while four chose to use remote pointing, and five switched between the two.

Display Distance. The distance to the display played a large role in participants' preferences, as well as in their decisions about which device to use in Phase II of the study. Because direct-touch requires a person to stand at close proximity to a large display, it can be difficult to obtain an overview of the entire canvas. For example, when using direct-touch in Phase I of the study, some participants reported that the lack of overview made it difficult to complete the task. One participant reported that standing in front of the display felt like "standing too close to a computer screen" and also reported that he was feeling heat emitted by the plasma display.

We observed several strategies to deal with this lack of overview. Many people stepped back and forth between trials to transfer from reading instructions to interacting with the display. Other people stood still and used wide head movements to read instructions. In Phase II, participants continued to use broad head movements when using direct touch to look up a desired operation on the menu in the screen's top left corner (we chose this default location to prevent people from "losing" the menu after stepping back and shifting their gaze). In contrast, in the remote-pointing condition, people were able to see the whole display without moving their head. Several participants reported this as a major benefit over the direct-touch technique.

The proximity to the display also introduces the difficulty of reach. We observed that many participants had to make sideways steps to reach remote areas of the screen in both phases of the study. When participants sat and used remote pointing, their movement was typically constrained to small arm and wrist movements.

Control Space. Both interaction techniques are also characterized by differences in their control spaces. While broad arm movements covering the entirety of the wall display were common for direct-touch interaction, small movements of the Wiimote achieved similar results. The broad movements of the direct-touch interaction paired with the direct physical contact allowed participants to achieve high precision in their actions and good control over the changes they applied. They reported that it "feels like a physical connection" and is "more accurate" and provides "more control." However, participants also mentioned that their arm got tired after a while due to the repeated arm movements that they used for accomplishing, in particular, the create and align tasks. In contrast, participants used small movements in the remote-pointing condition. These fairly small movements of the Wiimote resulted in big actions on the screen and, thus, induced larger errors in pointing and dragging/moving. Participants who started Phase I with remote pointing reported a noticeable gain in precision and control when they switched to direct touch, especially during the create and align tasks. Some noted that the direct technique "felt more accurate." On the other hand, participants also reported that the Wiimote "becomes extension of your hand" after a while and feels like an "extension of your body" or even that the remote pointing feels like playing a video game.

Participants had several strategies for dealing with the inaccuracy. Some people rested their forearm on the armrest or on their lap, and pointed with the wrist. This technique seemed to be more comfortable than pointing with the whole arm. Some participants locked their forearm to the side of their body and tilted their entire upper body to point across the screen. One participant even held her right arm (holding the Wiimote) over her left arm (holding the Nunchuk) in the create and align tasks, thus compensating for vertical and horizontal precision with separate arms. According to the participant, this arm-crossing was not tiresome. Some people switched between devices to gain the advantages of both techniques, creating broad layouts with remote pointing and working out fine detail via direct touch.

Forgiving Imprecision. We did not observe any main effects or interactions involving the block factor in either average distance nor coordination. We did, however, observe a speed improvement over time. These results suggest that our participants tended to sacrifice accuracy for speed. We also observed behaviour consistent with these results. For example, many participants seemed to be less and less careful to keep their actions constrained to the target area as the blocks progressed. Some participants would blatantly choose to not completely fill the target area or ignore when their movement line was not straight, despite our initial instruction to stay within the target boundary.

We suspect that this behaviour may be partly due to the fact that the painting application is very tolerant of inaccuracy. The application is forgiving both on the small scale and on the large. For example, when creating strokes, the exact location of each stroke is constrained to be within the area of influence, but the strokes are randomly distributed at each time step. Also, for any of the actions provided, an action by the user will affect any stroke whose centre is in the area of influence, and so parts of the stroke may rotate outside of this area. These small-scale inaccuracies may have encouraged participants to favour speed in the first phase. On the large scale, the application allows the creation of non-photorealistic images and, specifically, invites abstraction and expressiveness (for some example results from Phase II see Figure 12). Small errors or inaccuracies are, therefore, not noticeable or even desired as part of the artistic exploration process or as part of the intended effect. Alternatively, errors can also be corrected easily and without penalty by erasing or painting over previous strokes. Consequently, in the second phase, we observed that people tended to be satisfied with a final image that reflected their intended or unintended level of abstraction and expressiveness.

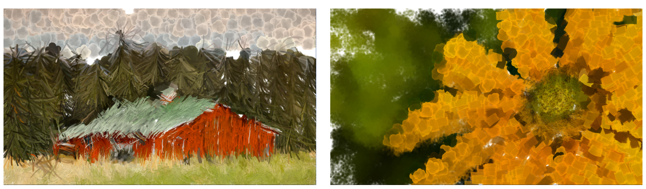

Figure 12: Two example results that participants created in Phase II within approximately 15 minutes.

We were initially surprised that many participants chose to use remote pointing in the second phase, despite its obvious performance drawbacks. However, because the application was forgiving, participants may have recognized that they could sacrifice accuracy to achieve speeds close to those of direct touch, and therefore leverage some of remote pointing's other benefits. Some participants also commented that using the Wiimote was more fun.

Handedness. We suspected initially that people would prefer to use the Nunchuk in their non-dominant hand and to touch or point with their dominant hand. Previous research has shown that the non-dominant hand is best suited to both actions that do not require precise movement and actions that use small motions with the thumb or wrist [11]. Because the Nunchuk interaction primarily required only a ballistic movement to select one of the eight corners on the joystick, and this motion was activated with the thumb, this mapping is consistent with this literature. However, in the direct-touch condition, seven participants chose to hold the Nunchuk in their dominant (right) hand and to interact with the display with their non-dominant (left) hand. Furthermore, this did not seem to adversely affect their performance. We suspect that this choice is again due to the forgiving nature of the application. Because the actions required by direct touch are not precise by nature and because the device offers more control than the Wiimote, participants may have decided that the Nunchuk interaction required their dominant-hand abilities. One of the participants who chose to interact in direct touch this way commented that he made this choice because he wanted more control of the Nunchuk menu.

CONCLUSION

We summarize our findings as follows:

- Direct touch was shown to be faster and more precise than remote pointing.
- With remote pointing, people are able to achieve speeds similar to direct touch by sacrificing accuracy.
- Applications that are tolerant of imprecision or invite exploration may alleviate some of the disadvantages of otherwise less efficient interaction methods.
- People had mixed preferences for remote pointing and direct touch. In general, we did not notice a correlation between preference and performance. Some preferred direct touch for its improved performance, but some preferred remote pointing for the ability to have an overview and for less fatiguing movement. Others preferred to switch between the two techniques to achieve different levels of precision and control at different times.

Both bimanual interaction techniques were shown to be suitable for the Interactive Canvas. We believe that the best solution is to allow both techniques of interaction. This redundancy allows people to choose the appropriate tool for the appropriate task. For example, when creating strokes to fill the canvas, a person can sit and view the entire screen at once and avoid the need to reach across the entire display, but when controlled motion is required to, e. g., align strokes, a person can stand and interact directly with the canvas.

In general, our new interaction techniques are a step toward providing more freedom to create and interact with non-photorealistic rendering. We believe that this form of redundant interaction is particularly useful in this domain, but would be beneficial in the design of other applications that require such freedoms.

FUTURE WORK

In the future, we would like to observe artists using our system over longer periods of time. The study presented in this paper typically lasted about an hour, and included many non-artists. We believe that it is particularly important to observe more long-term behaviour, since an hour-long session is not sufficient for people to begin to truly recognize the limitations and abilities of these techniques within this system. We would also like to observe children using these techniques with the Interactive Canvas, as many of our participants commented on their enjoyment level and on how it reminded them of paintings they did as children. One participant commented that "children would love this."

REFERENCES

1. Ahlström, B., Lenman, S., and Marmolin, T. Overcoming touchscreen user fatigue by workplace design. In *Poster Proc. CHI*, pp. 101–102, New York, 1992. ACM Press.
2. Baxter, B., Scheib, V., Lin, M. C., and Manocha, D. DAB: Interactive haptic painting with 3D virtual brushes. In *Proc. SIGGRAPH*, pp. 461–468, New York, 2001. ACM Press.
3. Buxton, W. and Myers, B. A. A study in two-handed input. In *Proc. CHI*, pp. 321–326, New York, 1986. ACM Press.
4. Card, S. K., English, W. K., and Burr, B. J. Evaluation of mouse, rate-controlled isometric joystick, step keys, and text keys for text selection on a CRT. *Ergonomics*, 21(8):601–613, Aug. 1978.
5. Forlines, C., Wigdor, D., Shen, C., and Balakrishnan, R. Direct-touch vs. mouse input for tabletop displays. In *Proc. CHI*, pp. 647–656, New York, 2007. ACM Press.
6. Gooch, B. and Gooch, A. A. *Non-Photorealistic Rendering*. A K Peters, Ltd., Natick, 2001.
7. Guiard, Y. Asymmetric division of labor in human skilled bimanual action: The kinematic chain as a model. *Journal of Motor Behavior*, 19(4):486–517, Dec. 1987.
8. Hertzmann, A. A survey of stroke-based rendering. *IEEE Computer Graphics and Applications*, 23(4):70–81, July/Aug. 2003.
9. Hinckley, K., Pausch, R., Proffitt, D., Patten, J., and Kassell, N. Cooperative bimanual action. In *Proc. CHI*, pp. 27–34, New York, 1997. ACM Press.
10. Isenberg, T., Miede, A., and Carpendale, S. A buffer framework for supporting responsive interaction in information visualization interfaces. In *Proc. C^5*, pp. 262–269, Los Alamitos, 2006. IEEE Computer Society.
11. Kabbash, P., MacKenzie, I. S., and Buxton, W. Human performance using computer input devices in the preferred and non-preferred hands. In *Proc. CHI*, pp. 474–481, New York, 1993. ACM Press.
12. Kalnins, R. D., Markosian, L., Meier, B. J., Kowalski, M. A., Lee, J. C., Davidson, P. L., Webb, M., Hughes, J. F., and Finkelstein, A. WYSIWYG NPR: Drawing strokes directly on 3D models. *j-TOG*, 21(3):755–762, July 2002.
13. Keefe, D. F., Acevedo Feliz, D., Moscovich, T., Laidlaw, D. H., and LaViola Jr., J. J. CavePainting: A fully immersive 3D artistic medium and interactive experience. In *Proc. I3D*, pp. 85–93, New York, 2001. ACM Press.
14. Lang, D., Findlater, L., and Shaver, M. CoolPaint: Direct interaction painting. In *Poster Proc. UIST*, 2003.
15. Lee, S., Olsen, S. C., and Gooch, B. Interactive 3D fluid jet painting. In *Proc. NPAR*, pp. 97–104, New York, 2006. ACM Press.
16. Myers, B. A., Bhatnagar, R., Nichols, J., Peck, C. H., Kong, D., Miller, R., and Long, A. C. Interacting at a distance: Measuring the performance of laser pointers and other devices. In *Proc. CHI*, pp. 33–40, New York, 2002. ACM Press.
17. Parker, J. K., Mandryk, R. L., and Inkpen, K. M. TractorBeam: Seamless integration of local and remote pointing for tabletop displays. In *Proc. GI*, pp. 33–40, Mississauga, ON, Canada, 2005. CHCCS.
18. Ryokai, K., Marti, S., and Ishii, H. I/O Brush: Drawing with everyday objects as ink. In *Proc. CHI*, pp. 303–310, New York, 2004. ACM Press.
19. Schkolne, S., Pruett, M., and Schröder, P. Surface drawing: Creating organic 3D shapes with the hand and tangible tools. In *Proc. CHI*, pp. 261–268, New York, 2001. ACM Press.
20. Schwarz, M., Isenberg, T., Mason, K., and Carpendale, S. Modeling with rendering primitives: An interactive non-photorealistic canvas. In *Proc. NPAR*, pp. 15–22, New York, 2007. ACM Press.
21. Sears, A. and Shneiderman, B. High precision touchscreens: Design strategies and comparisons with a mouse. *International Journal of Man-Machine Studies*, 34(4):593–613, Apr. 1991.
22. Sellen, A. J., Kurtenbach, G. P., and Buxton, W. A. S. The prevention of mode errors through sensory feedback. *Human-Computer Interaction*, 7(2):141–164, Apr. 1992.
23. Shirai, A., Geslin, E., and Richir, S. WiiMedia: Motion analysis methods and applications using a consumer video game controller. In *Proc. Sandbox*, pp. 133–140, New York, 2007. ACM Press.
24. Shugrina, M., Betke, M., and Collomosse, J. Empathic painting: Interactive stylization through observed emotional state. In *Proc. NPAR*, pp. 87–96, New York, 2006. ACM Press.
25. Strothotte, T. and Schlechtweg, S. *Non-Photorealistic Computer Graphics. Modeling, Animation, and Rendering*. Morgan Kaufmann Publishers, San Francisco, 2002.
26. Vogel, D. and Balakrishnan, R. Distant freehand pointing and clicking on very large, high resolution displays. In *Proc. UIST*, pp. 33–42, New York, 2005. ACM Press.
27. Zhai, S. and Milgram, P. Quantifying coordination in multiple DOF movement and its application to evaluating 6 DOF input devices. In *Proc. CHI*, pp. 320–327, New York, 1998. ACM Press.