Bi-level and Anti-aliased Rendering Methods for a Low-Resolution 2D Vibrotactile Array

Christoph W. Borst   Alp V. Asutay
Center for Advanced Computer Studies,
University of Louisiana at Lafayette, USA
E-mail: {cborst, ava4782}@cacs.louisiana.edu

Abstract

We investigate rendering methods for 2D tactile arrays. We present four rendering methods, two of which are anti-aliased methods to improve display quality and smoothness. We describe three experiments to evaluate the methods. The first experiment compares them based on subjective ratings and finds that anti-aliased methods improve perceived quality and smoothness. The second experiment investigates the potential for one approach to communicate information at a sub-tactor resolution, based on the shortest line segment allowing identification of its direction. The third experiment is a shape discrimination experiment and finds that only one of the approaches allows reasonable discrimination for the selected shapes.

1. Introduction

Research involving low-resolution 2D tactile arrays is becoming increasingly common. For example, arrays have been placed in vests and on the backs of seats to communicate direction or attentional cues. The arrays are often constructed using low-cost tactile elements (tactors) such as Piezo speakers or vibrating DC motors, which are useful for creating low-power tactile displays that are portable, lightweight, inexpensive, and simple to implement. Low resolution can result in perceivable spatial aliasing artifacts – e.g., discrete jumps in a moving stimulus rather than the feeling of smooth motion. Earlier, we compensated for this with a rudimentary anti-aliasing approach to provide smoother feedback for our own array, pictured in Figure 1 and discussed further in [1]. The vibrating DC motors used in this array (and by others) do not allow frequency and amplitude to be controlled independently, and neither parameter can be controlled precisely. This raises concerns that anti-aliasing may not produce good results, e.g., due to beats (low-frequency pulses) that occur when multiple tactors vibrate simultaneously at different frequencies.

Existing work with 1D arrays suggests effective resolution may be increased using tactor activation sequences that repeatedly pulse tactors to take advantage of a "saltation" illusion [2]. The technique has also been applied to 2D arrays, e.g., in [3]. We instead investigate rendering techniques that are related to techniques from graphical rendering and two of which might be described as "intensity interpolation" to provide a sense of smooth motion.

We begin by presenting two bi-level and two anti-aliased rendering methods for displaying traces of geometric shapes on a 2D array. We then present three experiments that evaluated the rendering methods. The first experiment compared the methods based on subjective evaluation by users, the second investigated the potential for an anti-aliasing method to present information at a sub-tactor resolution, and the third compared the methods based on a shape discrimination experiment.

Fig. 1. Palm-sized vibrotactile array

Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems
0-7695-2310-2/05 $20.00

2. Rendering Methods

2.1 Overview and Common Elements

We describe four rendering methods. Based on our earlier experiences with a DC motor array, shapes drawn as static tactile "images" appear difficult to understand, and so all methods represent a shape using a moving point that traces its outline. We assume its path is described by parametric equation $\mathbf{f}(t)$ that computes a 2D coordinate tracing the shape as t goes from 0 to 1. For example, a line segment with endpoints $\mathbf{P}_0$ and $\mathbf{P}_1$ is described by $\mathbf{f}_{\text{line}}(t) = \mathbf{P}_0 + (\mathbf{P}_1 - \mathbf{P}_0)t$, and a circle with center (c_x, c_y) and radius r by $\mathbf{f}_{\text{circle}}(t) = (c_x + r \cdot \cos(2\pi t),\ c_y + r \cdot \sin(2\pi t))$. Rendering involves controlling parameter t, either by repeatedly adding a small increment or based on a value from a real-time timer, as $\mathbf{f}(t)$ is mapped to the array. We note our "closest-motor" and "area sampling" methods are also applicable when we do not have an equation but instead base stimulus coordinate on some other value, such as the touchpad input coordinate in our interpersonal haptic communication system [1].

The four methods differ in the choice of activated tactors and in their use of intensity levels. Figures 2, 3, and 4 summarize the differences in activated tactors for rendering of a line segment and are discussed further in the following subsections.

The discussions assume that tactors are centered in unit square cells of an axis-aligned grid with its lower-left cell corner at the origin, and that tactor intensities range from 0 (off) to 1 (full intensity).

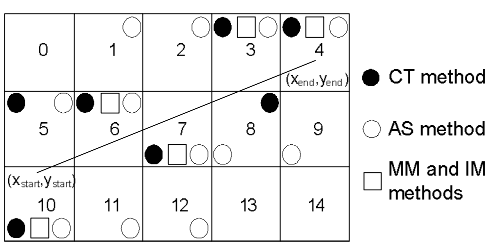

Fig. 2. Tactors activated by four methods

2.2. Closest-Tactor Method (CT)

The first method, called closest-tactor, is a bi-level method (tactor intensity is either 0 or 1). As the shape is traced, it activates the tactor with center closest to $\mathbf{P} = (p_x, p_y) = \mathbf{f}(t)$, and any other previously activated tactor is turned off. To identify the closest tactor, the method finds the grid cell containing $\mathbf{P}$. If each cell is addressed by the 2D index that is also the coordinate of its lower left corner, then the tactor to be activated is in cell (floor(p_x), floor(p_y)). If $\mathbf{P}$ is equidistant from multiple tactors, the upper/right tactor is chosen.

The resulting tactor activation sequence is seen in Figure 3, for the line segment of Figure 2 and with total tracing time of one second. Note that tactor activation durations are non-uniform.

The CT method is simple but produces significant spatial aliasing artifacts – rendered point motion occurs in discrete jumps rather than smoothly. Furthermore, the nonuniform durations might result in some tactor activations feeling "weak" compared to others, so there are concerns that the intensity or width of the line may appear to vary during the trace.

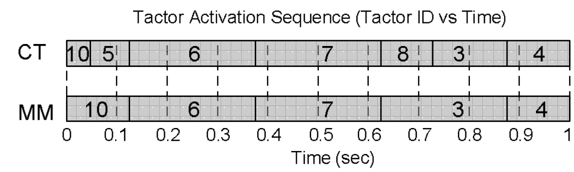

Fig. 3. Tactor activation sequences for the bi-level CT and MM rendering methods

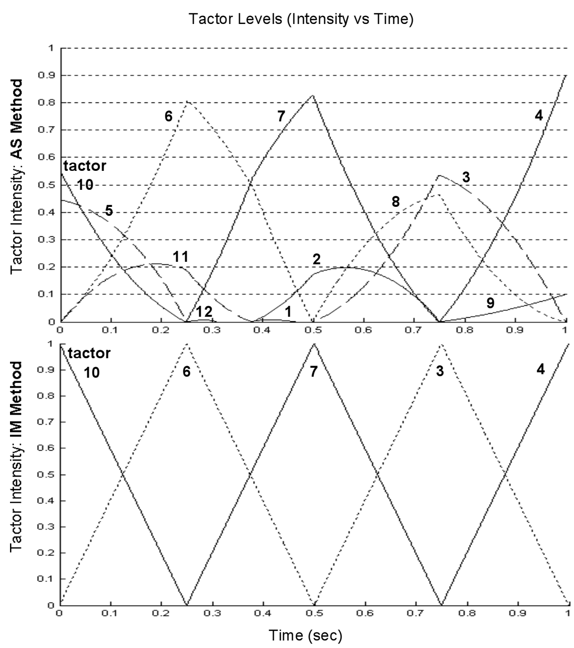

Fig. 4. Tactor activations for the anti-aliased AS and IM rendering methods

2.3. Midpoint Method (MM)

Our next method shares its name with an approach to graphical scan conversion and in fact generates the same sequence of elements (pixels/tactors). Consider a line segment with slope between -1 and 1 (other slopes are handled by symmetry). The MM method activates only one tactor per array column. Let mid_x be the x-coordinate of a column's center, and mid_y be a y-coordinate such that point (mid_x, mid_y) is on the line. The tactor centered at (mid_x, floor(mid_y)+0.5) is the closest tactor to the line in the column (or tied with the tactor below it). This tactor is activated when the trace $\mathbf{f}(t)$ is in the column and deactivated when it exits the column or ends. In practice, the sequence of tactors can be identified efficiently using basic graphical scan conversion routines [4]. A related midpoint technique exists for circles.

The resulting activation sequence is seen in Figure 3. Comparing to the CT method, we note that fewer tactors are activated and that tactor durations for line rendering with the MM method are uniform except for endpoint tactors. In fact, the MM method for lines can be implemented without repeated evaluation of the parametric equation – the total trace time is simply distributed between columns evenly, except at endpoints where it is assigned in proportion to the span of the line segment in the columns. Although the MM method suffers from the same basic spatial aliasing problem as the CT method, fewer jumps occur since fewer tactors are involved.

2.4. Area Sampling Method (AS)

If further tactor intensity control is available (i.e., it is not bi-level), it can be used to create smoother changes during rendering. One approach is to use area sampling methods as used to perform anti-aliased scan conversion in some graphics systems. We use an unweighted area sampling method that considers a unit square centered on the moving point **P**. At any point in time, this unit square overlaps one, two, or four array grid cells. Each tactor in an overlapped grid cell is activated with intensity equal to the area of intersection between its grid cell and the **P**-centered unit square, and any other previously activated tactors are turned off. To identify the (at most) four tactors activated, the coordinates of the **P**-centered box's corners are used to identify the four tactor grid cells containing them, each by an indexing operation as in Section 2.2. Per-cell overlaps are then computed using the known grid cell boundary coordinates and **P**-centered box boundary coordinates.

The resulting activation sequence is seen in Figure 4. There are no discrete jumps in tactor intensity. The activation curves consist of quadratic pieces that sum to one for any time point. Although discrete jumps are eliminated, this method raises concerns about rendered traces appearing thickened, since multiple tactors are activated at once and more tactors are activated in total. This motivates our fourth method and a gamma correction described later.

2.5. Interpolated-Midpoint Method (IM)

Our fourth method chooses the same tactors as the MM method but interpolates intensity. This minimizes the number of tactors while providing smooth changes, thereby sharing AS's ability to provide smooth sensations without concern about thickening of traces resulting from increased tactor count.

Again consider a line segment with a slope between -1 and 1, noting others are handled by symmetry. Consider point (x_t, y_t) = $\mathbf{f}(t)$ generated during a trace of the line using its parametric equation. The x-coordinate, x_t, lies between the tactor grid columns centered at x_k = [floor(x_t – 0.5) + 0.5] and x_{k+1} = (x_k + 1). Let one tactor per column be selected as in the MM method. The tactor with x-coordinate x_k is activated with intensity I_k = (x_{k+1} – x_t) and the tactor with x-coordinate x_{k+1} is activated with intensity (1.0 – I_k).

The resulting activation sequence is seen in Figure 4. As can be expected from the above expressions, the plot reflects linear interpolation of tactor intensities.

2.6. Threshold/Gamma Correction

We add "threshold" and "gamma" parameters to provide additional control over sensations for the AS and IM methods. Threshold, μ, sets a minimum tactor activation level. This can compensate for tactors that require significant nonzero voltage to respond. Gamma, γ, controls a gamma correction that is analogous to gamma correction for visual displays (used to compensate for nonlinearities in a monitor's response to electron gun levels). It provides some correction for nonlinearities in tactor response or perceptual nonlinearities, if known. Increasing the value of either parameter increases tactor intensities, so the parameters are useful for controlling perceived intensity of the rendering methods. Proper choice of values should allow AS and IM methods to produce perceived intensity similar to that of any one of the other methods. Parameters can also be adjusted based on users' subjective opinions about resulting quality.

Given nonzero tactor intensity I, the equation used to compute the new intensity $I_{adjusted}$ is:

$$I_{adjusted} = (1-\mu)(I)^{\frac{1}{\gamma}} + \mu$$

Figure 5 illustrates this function. In the left plot, μ is 0.33 as γ is varied from 0.3 to 7.0. In the right plot, γ is 0.7 while μ is varied from 0.0 to 0.6.

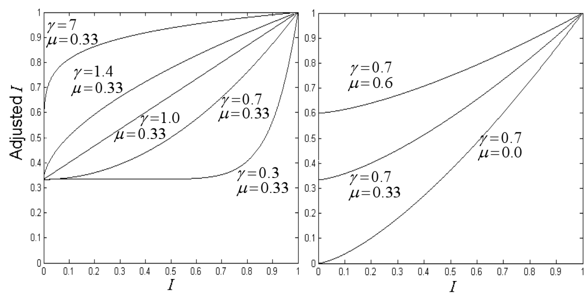

Fig. 5. Threshold and gamma correction

3. Experiment Methods

3.1. Experiment Design Overview

We evaluated the rendering methods using three experiments. The first experiment compared the four methods based on users' subjective ratings of quality, smoothness, and thickness. We used a within-subjects design with a randomized order for presenting methods. The independent variable was the rendering method used and the dependent variables were the resulting ratings. The research hypothesis was that the different methods produce different ratings, except in the case of thickness, because we intentionally set threshold and gamma parameters in an attempt to provide roughly equal perceived intensities.

The second experiment investigated the possibility that spatial anti-aliasing can allow users to perceive changes at a sub-tactor resolution. It considered only the AS method and used a tracking procedure to find the shortest line length for which subjects had a 70.7% success rate at identifying line direction based on 4-alternative questions. This was repeated for four different line types, in randomized order.

The third experiment was a within-subjects experiment to compare the four methods for shape discrimination between a circle and a square. The independent variable was the rendering method used and the dependent variable was the success rate. Methods were presented in randomized order. The research hypothesis was that different rendering methods would produce different success rates.

Except as noted in the Experiment 2 procedure, values μ=0.33 and γ=0.7 were used for all rendering performed by AS and IM methods. These were chosen to produce similar overall perceived rendering intensity for all methods while providing good sensations, based on the judgment of the authors.

Wherever orders were randomized, we mean order was assigned randomly from the permutations (ABCD, DCBA, ACBD, DBCA, BADC, CDAB, BDAC, CADB, ADBC, CBDA, BCAD, DACB) such that no two subjects experienced the same order.

3.2. Participants

Twelve male subjects participated. Age ranged from 19 to 35 years with median of 26 years. Subjects reported having no previous experience with tactile feedback devices beyond video game controllers and cell phones. Subjects were not compensated.

3.3. Materials

We used the 30-tactor palm-sized array shown in Figure 1 and detailed in [1, 5]. The tactors are low-cost vibrating DC motors in a 5 × 6 grid with inter-tactor spacing of about 18 mm (center-to-center). They are mounted to a project box using foam pads to help isolate them from the box and from each other. Pulse-width modulation (PWM) with a switching rate of roughly 300 Hz is used to provide 23 intensity levels. We measured the fundamental vibration frequency of the tactors to range from 27 Hz to 100 Hz, depending on the PWM pattern and the contact pressure.

Other equipment included a standard PC and liquid-filled headphones rated for 29db attenuation. The PC provided instructions, gathered responses, and controlled the array. Subjects wore the headphones for the duration of all experiments.

3.4. Procedure

3.4.1 Pre-experiment steps. Before the experiment, each subject reviewed and signed a consent document and answered a questionnaire requesting background information. The subject was then seated at the PC and the array and donned the headphones.

3.4.2 Experiment 1. Experiment 1 had three stages. The first was a preparation stage. The subject placed a hand on the array and freely selected between four rendering methods by pressing four number keys while sensing a line segment moving across the array from left-to-right with slope 0.58 (method assignment to number keys was randomized). Tracing speed was set for a duration of 2 seconds per trace. The subject pressed a letter key when ready to provide ratings.

0-7695-2310-2/05 $20.00 © 2005 **IEEE**

In the second stage, subjects rated line segment quality, smoothness, and thickness for each of the four methods. Quality was rated first, followed by smoothness, followed by thickness. Order of methods was randomized. Each rating was provided on a scale from 0 to 100 by keyboard text entry by the free hand while the other hand remained on the array and received the relevant feedback (the line was traced repeatedly using a 2-second trace duration with a half-second pause between traces).

The third stage repeated the second except that the shape was a circle centered on the array with radius of 2 grid units (36 mm), traced counter-clockwise.

3.4.3 Experiment 2. The subject removed the hand from the array and rested briefly before Experiment 2. Experiment 2 had four stages, each being a similar 30-question tracking experiment using 4-alternative questions, with a different stimulus type for each of the four stages. Only the AS method was used for rendering. Each question required the user to indicate the direction of a short line segment being repeatedly traced on the array, by pressing an arrow key to indicate up, down, left, or right (actual line direction was randomly selected to be one of these). Software provided text feedback indicating if answers were correct or not, also indicating the correct answer whenever the subject was incorrect. Line length for each stage began at 2.0 grid units (36 mm). It was reduced by 31.5% for two successive correct answers and increased by 46% for each incorrect answer. Based on Levitt [6], the procedure converges to a line length that results in about 70% correct responses.

Lines were traced repeatedly with 2-second trace durations and a half-second pause between traces. For one stage, lines had a tactor center as their midpoint and were rendered with parameters μ=0.33 and γ=0.7. Another stage rendered lines with the midpoint being the corner point of a grid cell (a point equidistant from the four nearest tactors) with the same parameter values. The other two stages rendered these two line segments with μ=0.0 and γ=1.4. Order of stages was randomized. Between successive stages, the subject removed the hand from the array and rested briefly.

3.4.4 Experiment 3. The subject removed the hand from the array and rested briefly before Experiment 3. For each of the four rendering methods, the subject was then asked to identify 10 rendered shapes, each being either a circle or a square. So, there were four sets of 10 discrimination questions. The order of the four sets was randomized. For each set, five of the drawn shapes were circles and five were squares, presented in random order. Squares and circles were centered on the center array column and on the fourth row. Squares had an edge length of 4 grid units (72 mm) and circles had a diameter of just over 4 units. Software provided text feedback indicating if answers were correct or not. Shapes were retraced continuously in a counterclockwise direction with a per-trace duration of two seconds. A total time limit of 20 seconds was given for each question.

4. Experiment Results and Discussion

4.1 Experiment 1

For Experiment 1, we computed quality, smoothness, and thickness scores, each by summing corresponding ratings for lines and circles (e.g., quality is the sum of line quality and circle quality ratings). Figure 6 shows resulting scores with a box plot for each of the four rendering methods and each of the three scores. Boxes show medians and quartiles, with whiskers showing the full range of scores except for indicated mild and extreme outliers.

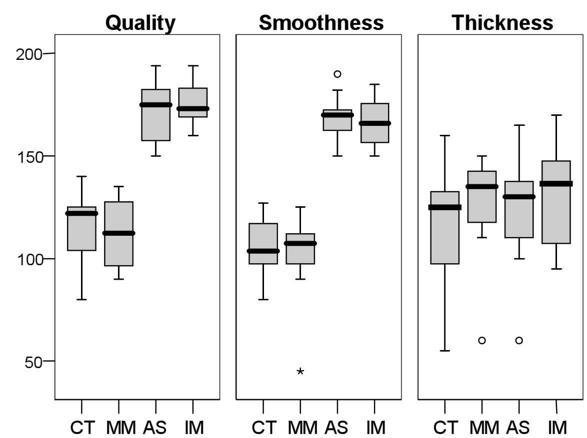

Fig. 6. Experiment 1 quality, smoothness, and thickness scores for the four methods

For each scored property, we compared the methods using bonferroni-adjusted pairwise Wilcoxon signed-rank tests for related samples. The resulting p-values are given in Table 1.

The results show that both anti-aliased methods result in significantly better perceived quality and smoothness than either bi-level method. No significant differences were detected between the two anti-aliased methods or between the two bi-level methods. No significant differences were detected based on

Proceedings of the First Joint Eurohaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems
0-7695-2310-2/05 $20.00 © 2005 **IEEE**

thickness (recall that this was expected due to the particular choice of gamma and threshold values).

Experiment 1 confirms that anti-aliasing improves subjective quality and smoothness of feedback.

Table 1. Experiment 1 pairwise Wilcoxon p-values, after multiplication by 6 for Bonferroni adjustment. Values below 0.05 are significant.

	CT, MM	CT, AS	CT, IM	MM, AS	MM, IM	AS, IM
Quality	1.000	0.012	0.012	0.012	0.012	1.000
Smoothness	1.000	0.012	0.012	0.012	0.012	1.000
Thickness	0.168	1.000	0.186	1.000	1.000	1.000

4.2 Experiment 2

Experiment 2 estimated the shortest line length for 70.7% correct response for line direction (61% correct sensing of direction, 9.7% correct guesses, and 29.3% incorrect guesses). To estimate line length from the tracking procedure history, we average lengths from the last four trials that produced changes in direction (we average the last two peaks and last two valleys).

Results are shown in Figure 7. Recall only the AS method was used. The four conditions were:

L1 – midpoint at cell corner, μ=0.33, γ=0.7
L2 – midpoint at cell corner, μ=0.0, γ=1.4
L3 – midpoint at tactor center, μ=0.33, γ=0.7
L4 – midpoint at tactor center, μ=0.0, γ=1.4

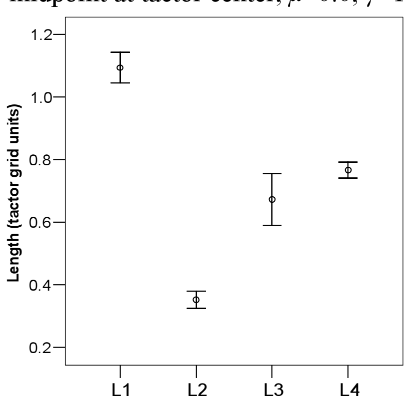

Fig. 7. Mean line lengths (+/- 1 SE) for 70.7% correct response for conditions L1, L2, L3, L4

To understand the results, we also consider what would be expected with a bi-level method. A line with midpoint at a tactor center is a worst-case for the bi-level methods in that any line shorter than 1 grid unit activates only a single tactor, so no motion or direction could be perceived. So, L3 and L4 results show that the AS method can communicate movements in this region not possible with the bi-level methods. There was large variability for L3, with one-quarter of subjects converging at lengths below 0.4 units (7.2 mm), well below the mean. We explain the ability of some users to detect short L3 lines and argue that shorter lines can be made detectable as follows: Three tactors were involved – a middle tactor behaving identically for all directions and two adjacent tactors on opposite sides. As the trace passed through the middle tactor's center, the adjacent tactor on one side deactivated as the opposite tactor activated. Increased μ makes the transition more noticeable by increasing the level of change at the moment of crossing, and in the extreme case this becomes a sudden large jump in stimulus. Parameters μ and γ can be set to make arbitrarily short line segments behave like longer ones.

Conversely, a line with midpoint at a cell corner is a best case for the bi-level methods in that even extremely short lines cross a cell boundary and result in a sudden discrete jump, revealing direction. The change is less dramatic with anti-aliased methods producing smoother changes. The L2 result shows that a line much shorter than one grid unit is nonetheless perceivable for certain parameter values, while the L1 result shows this was not accomplished with the other investigated values. Here, four tactors are activated for a line shorter than 1 grid unit, and all are active with equal intensity when the trace crosses the cell corner (before μ and γ adjustment, this intensity is 0.25). To emphasize changes in this region, parameters should be set such that small movements produce noticeable changes in actual tactor response. Based on frequency and amplitude response plotted in [5], tactor response for our device has highest slope for intensity of roughly 0.2–0.4 (0.2 is the smallest intensity for any response). The L1 condition used a threshold of μ=0.33, which, applied to the pre-corrected intensity of 0.25 for the trace at the cell corner, did not make good use of tactor response. On the other hand, L2 made good use of the response characteristic.

Overall, Experiment 2 illustrates tradeoffs in setting μ and γ for the perception of small movements and provides a starting point for further work to find values producing a consistent result throughout the array. Based on the experiment, we expect to find values that allow sub-tactor line lengths (movements below 1 grid unit) to communicate information throughout the array area. L2 and L3 discussions showed how changes can be amplified for specific regions, but the L1 result illustrates how amplifying changes in one region can compromise sensations in another. L2 and L4 results show a single choice of parameters supporting sub-tactor line lengths for both types of regions, which is not possible with the bi-level approaches due to the array's physical resolution. At least in this sense, we say the AS method can communicate information at a sub-tactor resolution.

4.3 Experiment 3

Experiment 3 discrimination scores were recorded for each subject on a per-method basis, with score being the number of shapes identified correctly. An average score of 5 is expected if no discrimination was possible (i.e., if the subject was just guessing). The resulting scores are shown in Figure 8.

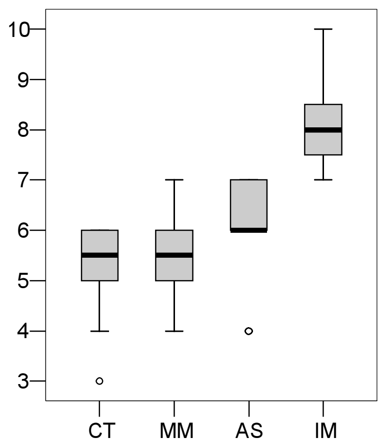

Fig. 8. Experiment 3 discrimination scores for four rendering methods CT, MM, AS, IM

Table 2. Experiment 3 pairwise Wilcoxon p-values, after multiplication by 6 for Bonferroni adjustment. Values below 0.05 are significant.

CT, MM	CT, AS	CT, IM	MM, AS	MM, IM	AS, IM
1.00	0.276	0.012	0.672	0.012	0.018

We compared methods using bonferroni-adjusted pairwise Wilcoxon signed-rank tests for related samples. Table 2 shows resulting p-values. Subjects were shown to be significantly better at discrimination with the IM method than with any other method. No other significant effects were detected.

Based on plotted results and subject comments, discrimination was difficult with our palm array. Differences between shapes were only detected reasonably with the IM method. We were surprised by the difference between MM and IM, since both activate the same tactors. One possible explanation is that the IM interpolation causes circle regions near array corners to appear less intense and less likely to be confused with square corners. However, inspection of MM data revealed that 62% of squares were misidentified compare to only 25% of circles, suggesting subjects were more likely to perceive squares as circles than circles as squares. Good discrimination may require a technique such as reparameterizing traces to emphasize distinct features of shapes, e.g., to spend more time in square corners.

5. Conclusion

Four rendering methods for 2D tactile arrays were presented and evaluated. We showed that anti-aliased methods significantly increase perceived quality and smoothness. We demonstrated that an anti-aliased method allows line direction to be identified for sub-tactor line lengths and illustrated how this relates to intensity correction parameters and to the region in which the stimulus occurs. Shape discrimination was difficult with our array, but one method allowed better discrimination than the others for the chosen test case.

We used low-cost DC motors that do not allow precise control and noted a possible problem with beats (in the Introduction). Based on the quality and smoothness ratings, this does not preclude overall improved display quality with anti-aliased approaches. Anecdotally, perceived beats are minimal since tactor frequencies change during display, so the low-frequency pulses are transient. An investigation of rendering methods using higher-quality tactors may reveal additional differences between methods.

References

[1] C. W. Borst and C. D. Cavanaugh, "Touchpad-Driven Haptic Communication using a Palm-Sized Vibrotactile Array with an Open-Hardware Controller Design," EuroHaptics conference, Munich, 2004.

[2] R. W. Cholewiak and A. A. Collins, "The generation of vibrotactile patterns on a linear array: Influences of body site, time, and presentation mode," *Perception and Psychophysics*, vol. 62, pp. 1220-1235, 2000.

[3] H. Z. Tan, R. Gray, J. J. Young, and R. Traylor, "A Haptic Back Display for Attentional and Directional Cueing," *Haptics-e, www.haptics-e.org*, vol. 3, 2003.

[4] J. D. Foley, A. V. Dam, S. K. Feiner, and J. F. Hughes, *Computer Graphics: Principles and Practice in C*. Boston: Addison Wesley, 1995.

[5] C. W. Borst and C. D. Cavanaugh, "Haptic Controller Design and Palm-sized Vibrotactile Array," Technical Report, Center for Advanced Computer Studies, University of Louisiana at Lafayette, 2004.

[6] H. Levitt, "Transformed Up-Down Methods in Psychoacoustics," *The Journal of the Acoustical Society of America*, vol. 49, pp. 467-477, 1971.