Evaluation of Drawing on 3D Surfaces with Haptics

Youngung Shon and Sara McMains
University of California at Berkeley

Experiments demonstrate that stereo displays, higher simulated friction, and interpolated-normal force shading, unlike shadows, improve performance when drawing on 3D surfaces in haptic environments.

In this article, we describe experiments for evaluating the speed and accuracy of a haptic drawing interface, particularly in terms of the effects of changing two haptic characteristics—the magnitude of the friction force and the use of smooth, interpolated-normal force shading—in combination with the effect of stereo displays and shadows. Because surface geometry affects the sense of touch, we studied the efficacy of the haptic interface in the context of varying geometric characteristics of the surface drawing task.

Our experiments were motivated by the desire to find optimal settings for a collaborative haptic mark-up environment (CHaMUE) we are developing that allows remote users to draw geometric annotations directly on a part surface. Drawing is an indispensable tool for communicating ideas between designers, engineers, and end users, especially during the conceptual design phase of new product development. For designing 3D artifacts, however, 2D drawings that show only a single viewpoint aren't always the most intuitive way to communicate a design or mark it up with suggested geometric changes. With the help of VR equipment, it becomes possible to design, discuss, and mark up such artifacts in a more realistic 3D environment using haptic input devices.

Giving users adequate feedback while they perform 3D drawing on a computer is much more challenging than with 2D drawing environments, however. *Proprioception*—sensation of limb position—provides only coarse information to users about the position of their drawing hands, even in the best-case scenario when the motion of a 3D input device is mapped to the virtual world without scaling or translation. Furthermore, a standard computer monitor shows only a single 2D projection of the virtual world at a time, and even with stereo displays, almost all users misjudge distances.[1] Therefore, in addition to stereo displays, our research tested the efficacy of showing the shadow of the drawing tool on the surface of the object being drawn upon.

Other research groups have developed systems for haptic drawing or painting on 3D surfaces.[2,3] SensAble Technologies sells FreeForm (http://www.sensable.com), a haptic design software package within which users can draw curves, including on the surface of existing geometry, that are used as input to geometry modification operators.

Other research experiments have found haptic interfaces helpful for targeting and training tasks, but haven't studied the effect of haptic parameter settings on 3D surface drawing.[4,5] The visual variables we consider have more often been studied outside of a haptics context.[6,7] One experiment found that binocular stereo improved users' targeting performance by a statistically significant amount and that shadows further enhanced the beneficial effect of stereo displays.[7] However, in an experiment with haptics, another research team reported that stereo displays did not improve performance in path-following tasks.[8] This last experiment, much like ours, had users trace paths on haptically rendered surfaces, but studied path tracing only on planar smooth surfaces.

Experimental setup

Eleven right-handed students participated in this experiment. None of the subjects had prior experience in VR environments with haptic devices. We conducted the tests with a Phantom Premium 3.0L haptic arm with a stylus end effector, a Fakespace ImmersaDesk, and CrystalEyes VR stereo shutter glasses with a Logitech ultrasonic head-tracking system. The setup is shown in Figure 1. The Phantom has a 6-DOF positional sensing input and provides 3-DOF force-feedback output, simulating single-point contact without torque.

The ImmersaDesk supports stereo with a 120-Hz refresh rate and no visible ghosting or crosstalk between the images for the left and right eyes. These full-screen images are alternately displayed each refresh cycle while signals from an emitter for the stereo

Published by the IEEE Computer Society

glasses synchronize shuttering the lens in front of the opposite eye to block its view of the screen temporarily. We adjusted the camera positions used to generate the images for each eye according to the user's head position relative to the screen, as reported by the ultrasonic tracking system.

We informed the subjects that both their speed and accuracy would be measured as they used the Phantom to draw over curves on the surfaces of various 3D virtual objects. We displayed the 60 cm target curves as 5-pixel-wide red lines on white objects measuring approximately 20 × 20 × 20 cm. We positioned the virtual objects at a fixed position, but the head-tracking feature provided a natural and intuitive way to change the viewpoint. The target curves were generally not completely visible from the initial viewpoint, but all of the curves' initially hidden portions could be made visible with small adjustments of head position.

To draw, users would press a button on the Phantom stylus while in contact with the virtual object. Doing so would create a 5-pixel-wide blue line, overwriting any portions of the target line it covered. The virtual drawing tool was displayed as a paintbrush; when in contact with the object, a small red dot slightly larger than the line width of the target curve was displayed to indicate the exact location where the user line would be drawn. Without any prompting from the experimenters, most subjects spent 2 to 3 seconds haptically exploring each new surface by touching it with the brush before they pressed the button to start drawing on the curves.

To transition between trials, subjects pressed a virtual button labeled "load next" when they felt that the curve that they had drawn was complete. After every 40 trials, subjects took a 2-minute mandatory break. In addition, subjects took voluntary breaks at any time by pressing the virtual button labeled "break" instead of the "load next" button. During these breaks, subjects were not able to see anything but a "resume" button and the virtual drawing tool.

Because the subjects had no experience using haptic devices prior to this experiment, their performance often improved with practice, with varying amounts of time for the performance of different subjects to stabilize. We analyzed accuracy and speed during practice trial sets and gave additional practice sets to subjects whose performance did not fall within the 95 percent confidence interval calculated on the basis of the previous trials recorded.

Experimental variables

In our experiment we manipulated eight variables: two visual variables, two haptic variables, and four geometric variables. Table 1 shows a summary of these variables, their abbreviations, and the two levels tested for each.

1 Student using the haptic drawing system.

Table 1. Independent variables and their tested levels.

Factor	Level +1	Level −1
Stereo (S)	On: eye separation = 63 mm	Off: eye separation = 0 mm
Shadows (W)	On	Off
Normal shading (N)	On: interpolating	Off: flat
Friction (T)	$\mu_s = 0.5$ $\mu_d = 0.4$	$\mu_s = 0.1$ $\mu_d = 0.08$
Convexity (C)	Convex surface	Concave surface
Smoothness (D)	Faceted surface	Smooth surface
Curve direction (G)	Curve along gradient	Curve along isoheight
Hill frequency (F)	12 hills on surface	4 hills on surface

Stereo display

The first visual variable is the presence or absence of stereo viewing. We used the head-tracking system to set the viewpoint under both conditions and generated stereo images with two virtual cameras controlled by the head position relative to the screen, separated by the eye-separation value—otherwise known as the interpupillary distance. For nonstereo conditions, the eye separation is simply set to zero, effectively placing both cameras at the average eye position and thus eliminating parallax. Subjects wore the stereo shutter glasses even under nonstereo conditions so we could evaluate the effect of stereo independently from the shuttering effect, which darkens the perceived image.

We placed the virtual drawing tool so that it appeared to extend from the physical position of the Phantom stylus in the stereo mode, calibrated so that the discrepancy between the virtual and physical positions was less than 1 inch.

Drawing tool shadow

The second visual variable is the presence or absence of a cast shadow of the virtual drawing tool on the virtual object. Cast shadows can act as a visual cue about the distance between two objects. With our virtual paintbrush, the distance from the brush tip to the surface to be drawn upon can be estimated by looking at the distance from the brush tip to its shadow on the surface (see Figure 2, next page). We implemented real-time shadows using the hardware-accelerated volumetric-

2 **Shadows help provide a visual cue that the paintbrush on the left is touching the surface, while the paintbrush on the right is not.**

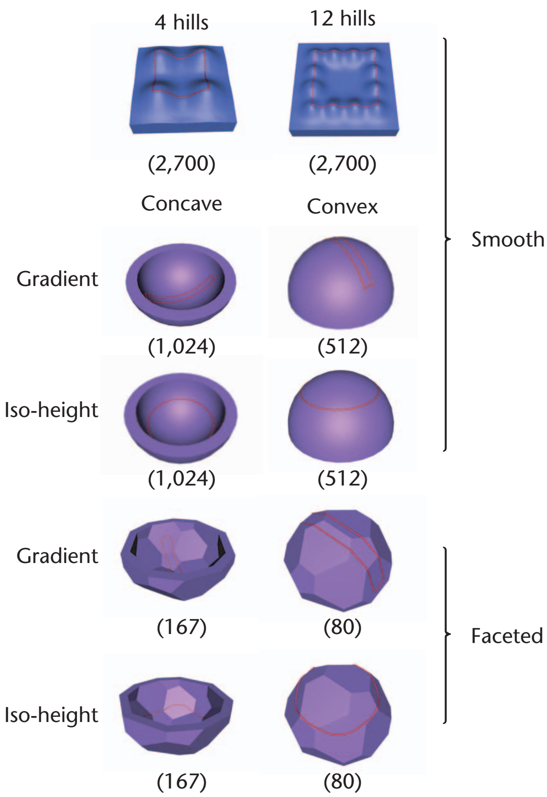

3 **Combinations of geometric factors. The *z* direction is rotated slightly out of the projection plane in the view pictured. The numbers in parentheses indicate the number of faces of each 3D object.**

shadow method.[9] Even for nonshadow conditions, we performed the additional rendering pass and calculations for the shadows so both conditions would have the same latency introduced by shadow computations.

Interpolated-normal force shading

The first haptic variable is the presence or absence of smooth, interpolated-normal force shading. The direction of the force feedback felt by a user when the virtual paintbrush is in contact with an object is affected by the surface-normal vector direction. Without interpolated-normal force shading, surface-normal vectors are invariant over a triangle that contains a contact point anywhere on its surface, analogous to flat-shading in computer graphics. In the interpolated-normal force shading mode, the surface-normal vector used for calculating force feedback instead interpolates three different normal vectors assigned to the vertices of the triangle.

This interpolated-normal force shading results in a smooth transition of the force direction when the trajectory of the contact point passes across one triangle to an adjacent triangle, making sharp edges feel smoother to users.[10] The vertex normals are the same normals used to calculate the surface-shading values interpolated for Gouraud shading. In our experiments, we always matched the display-shading mode to the force-shading mode. Users saw flat-shaded surfaces with noninterpolated haptic normals but saw smooth, Gouraud-shaded surfaces with interpolated-normal force shading for haptics.

The target curves did not appear any smoother with Gouraud shading, just as object silhouettes are not smoothed by Gouraud shading. Thus we could not make the visual display entirely consistent between the curves and surfaces in the interpolated-normal force shading mode.

Friction force

The second haptic variable is the amount of friction force used. We simulated friction forces by adding forces in the opposite direction from the current velocity when in contact with an object. We tested a low-friction condition using $\mu_s = 0.1$ as the static-friction coefficient and $\mu_d = 0.08$ as the dynamic-friction coefficient, and a high-friction condition using $\mu_s = 0.5$ and $\mu_d = 0.4$. We chose these coefficients to approximate metal on metal under well-lubricated conditions and rubber on metal under poorly lubricated conditions.

Convexity of objects

We tested combinations of four geometric variables, illustrated in Figure 3. The first variable was whether curves were on convex or concave surfaces. For the convex condition, subjects drew on the curved surface of a hemisphere. For the concave condition, subjects drew on the inside surface of a bowl.

Object smoothness

The second geometric variable is the smoothness or discontinuity of the object surfaces under the target curves. Under both conditions our virtual test objects were ultimately composed of triangular facets. But both visually and haptically, users perceive the finely faceted objects as smooth even without interpolating surface normals. Due to the relatively low complexity of even the finely tessellated 3D objects (maximum 2,700 facets), the system maintained stable update rates of 29 to 30 Hz for graphical displays and 1,000 Hz for haptic displays.

Curve direction

The third geometric variable is the dominant direction, relative to the surface, of the curve to be traced. In this experiment, target curves ran either primarily along the surface gradient direction with respect to z (height) or primarily along the isoheight direction, with the z direction perpendicular to the screen. These directions depend on the object's current orientation, but because we didn't allow the subjects to rotate the objects, the dominant curve direction is maintained.

Hill frequency

The nature of repetitive undulations of the object and target curves is the final geometric condition: comparing an object and target curve to ones with with lower curvature and less frequent changes in sign of curvature. The undulating curves follow hills on the surface. As seen at the top of Figure 3, we used one surface that has 4 hills and another surface that has 12 hills.

We did not test the hill frequency in combination with the other geometric variables. While users trace the undulations of the consecutive hills, they alternate between convex peaks and concave valleys; therefore, the surface is both convex and concave. For our target curve to trace over all the hills' peaks, its direction is always along a maximum gradient direction with respect to z. This means that the curve direction cannot be combined with hill frequency. Although hill frequency could conceivably be combined with the smoothness variable, it's not clear whether discontinuous hills of different sizes should keep the angle of the discontinuities constant or keep the number of discontinuities constant. We therefore decided to test hill frequency separately from the other geometric variables.

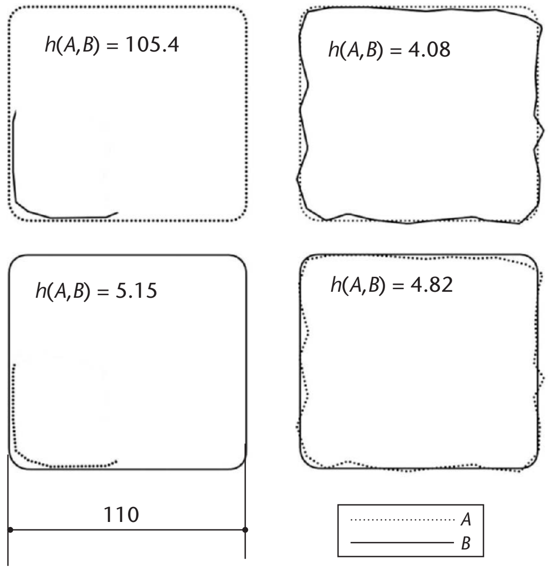

4 The choice of assignment of curves to *A* or *B* in the *h*(*A*,*B*) calculation can significantly affect the result. Here the set represented by the solid lines has an infinite number of points.

Dependent variables

The observed variables we measured included the speed and accuracy of the user-drawn curves. We recorded the total time—drawing and nondrawing—that subjects spent for each trial. Timing started when the 3D object loaded and stopped when the subject pressed the "load next" button. We measured both the average and worst-case error distance between the target curves and the user-drawn curves. We used the Hausdorff distance between two point sets A and B:

$$h(A, B) = max_{a \in A}\{min_{b \in B}\{distance(a, b)\}\}$$

This equation measures the worst-case error. For the distance function, we used the geodesic distance along the surface using the fast marching method.[11]

Because the Hausdorff distance is not symmetric—$h(A,B) \neq h(B,A)$—we had to consider how to assign the target curves and user-drawn curves to the two sets. Figure 4 shows how the wrong choice of assignment can lead to misleading worst-case error measurements. Because we wanted a higher error for an incompletely drawn user curve, we assigned the user curve points to set B. For set A, we parameterized the target curves by arc length and uniformly sampled 500 points.

However, we found that the worst-case error captured by the Hausdorff distance could vary by large amounts. The outliers dominate the statistical analysis, making it difficult to draw meaningful conclusions from the data. Therefore, we used the following variant:

$$h_{50}(A,B) = \frac{1}{50}\sum_{i=1}^{50} i^{th}\max_{a \in A}\left\{\min_{b \in B}\left\{\text{distance}(a,b)\right\}\right\}$$

This equation determines the average of the first 50 maximum distances—in other words the worst 10 percent of errors for each trial—as an error metric that compromises between average and maximum error.

Fractional factorial design

If we had tested all level combinations of our 8 two-level factors (a full factorial experimental design), the total number of trials would have been $2^8 = 256$. To increase the probability that we could draw statistically significant conclusions from our data, we wanted to replicate each combination we tested. Thus, with an average trial length of 25 seconds, it would have taken our subjects over 3.5 hours to complete a full factorial experiment with two replications, not including time for instruction, practice, and breaks.

Recruiting subjects for such a lengthy, repetitive experiment was not realistic, so we used a fractional factorial design, testing only a carefully chosen subset of combinations of levels. Measuring the magnitude of all possible combined interactions between any number of factors is not possible with such a fractional factorial design, but we can still draw meaningful conclusions on the basis of the reasonable assumption that higher-order interactions are likely to be negligible. For example, if analysis indicates that a particular performance improvement could be due to the presence of factor A or alternately to the combined presence of factors B, C, D, and E not attributable to B, C, D, nor E individually nor in pairs, then Occam's Razor suggests that the improvement is more likely to be due to factor A alone.

We can calculate several different effects, provided that all pairs of columns in the treatment-setting table are orthogonal to each other. We used orthogonal arrays to assign levels to factors.[12] With these arrays, all the main single-factor effects and the two-factor interaction effects we examined are confounded with higher-order interactions between three or more factors. Under the

Table 2. Examined and excluded two-factor interactions. (X indicates examined.)

Variables	S	W	N	T	C	D	G	F
Stereo (S)	-	X	-	-	X	-	X	X
Shadow (W)	X	-	-	-	-	-	-	-
Normal force shading (N)	-	-	-	X	X	X	X	X
Friction (T)	-	-	X	-	X	X	X	X
Convexity (C)	X	-	X	X	-	-	-	-
Smoothness (D)	-	-	X	X	-	-	-	-
Curve direction (G)	X	-	X	X	-	-	-	-
Hill frequency (F)	X	-	X	X	-	-	-	-

Table 3. Analysis-of-variance probability of null hypothesis. Statistically significant (low probability) entries are in bold.

Source	*p* (Average Error)	*p* (h_{50} error)	*p* (Time)
Stereo (S)	**0.0850**	**0.0185**	0.0405
Shadow (W)	0.5006	0.5554	0.8137
Force shading (N)	0.4865	0.5617	**0.0110**
Friction (T)	**0.0006**	**0.0055**	0.4456
Convexity (C)	**0.0003**	**0.0246**	0.1008
Smoothness (D)	**0.0000**	**0.0000**	**0.0000**
Curve direction (G)	**0.0982**	**0.0690**	**0.0800**
Hill frequency (F)	0.7856	0.5643	**0.0000**
S*W	0.2640	0.1888	0.1245
S*C	**0.0398**	**0.0432**	0.7323
S*G	0.3590	0.6641	0.6453
N*T	0.6485	0.6937	0.9141
N*C	0.2464	0.3309	0.4098
N*D	0.3588	0.1425	0.9083
N*G	0.3409	0.3284	0.5240
T*C	0.5289	0.7187	0.3988
T*D	**0.0904**	**0.0832**	0.5905
T*G	0.9856	0.7988	0.8199
S*F	0.7856	0.7231	0.4236
N*F	0.9223	0.8734	0.6958
T*F	0.6745	0.8238	0.3701

assumption that these higher-order effects are negligible, we can conclude that the variance we calculate is almost entirely due to the lower-order effects.

To reduce testing time, we studied only the interactions we suspected were more likely to be significant. We started with the assumption that interactions between three or more factors are negligible, reducing our choices to the 8 individual factors and 28 possible two-factor interactions. Even so, we didn't have enough time to explore all possible combinations of two factors; hence we needed to select suspected interactions to allocate our orthogonal array columns accordingly.

Among 28 possible two-factor interactions, we sacrificed measuring the following combinations that appeared less likely to be statistically significant:

- interactions between the presence of shadows and the geometric factors;
- visual cues (stereo and shadows) combined with haptic parameters, which are not visually displayed;
- stereo with smoothness, because smoothness isn't associated with depth; and
- interactions between geometric factors, because these factors are not variables that can be set to their optimal experimentally determined levels and combinations in real applications (because users will be able to draw wherever they want on any shape of object).

Table 2 shows interactions we chose to study and those we excluded. Just because we excluded an interaction from the study does not mean that its significance is necessarily small.

Because factor F (hill frequency) cannot be tested with different levels of the other geometric factors (C, D, and G), we broke the experiment into two subsets for the experimental design and statistical analysis. In the first subset, only C, D, and G appear as geometric factors. In the second subset, only F appears as a geometric factor. We used orthogonal arrays for the design and analysis of both subsets.

We were able to perform two replications of the experiment within 1.5 hours, including time for instruction and practice. We randomized the order of the 40 trials for each subject within each replication, with the exception of the stereo condition. During pilot studies, we found that transitions from the nonstereo setting to the stereo setting caused eye strain for some users. Therefore, we grouped the randomly ordered trials for each replication into nonstereo and stereo subsets, and used an ABBA design for these subsets to minimize the stereo transitions.

Results and analysis

We used analysis of variance (ANOVA) to analyze the data from the fractional factorial experiment. We performed statistical tests called *F* tests to test the null hypothesis for each factor—in other words to test the hypothesis that the factor had no effect. An *F* test calculates the probability *p* that the measured differences occurred by chance. We only reported results where the *F* test indicated a greater than 90 percent probability of significance, corresponding to criteria of $p < 0.1$. Table 3 summarizes the results of the *F* tests. For statistically significant effects, we also calculated the magnitude of the effect. We performed the statistical analysis in Mathworks Matlab and crosschecked the results with SAS Institute's Statistical Analysis Software.

The mean for the average error was 8.2 mm. The mean for h_{50} worst-case error was 19.3 mm. Five of the single factors had potentially statistically significant effects ($p < 0.1$) on both the error measures: stereo, friction, convexity, smoothness, and curve direction. Figure 5 shows the single-factor effects on average error. Stereo displays lowered the average error by 11 percent; in other words, stereo displays improved drawing accuracy.

Another statistically significant single factor is the friction-force condition. Subjects drew more accurately under the high-friction condition than under the low-friction condition. The average error was reduced by 1.5 mm (a 15 percent reduction).

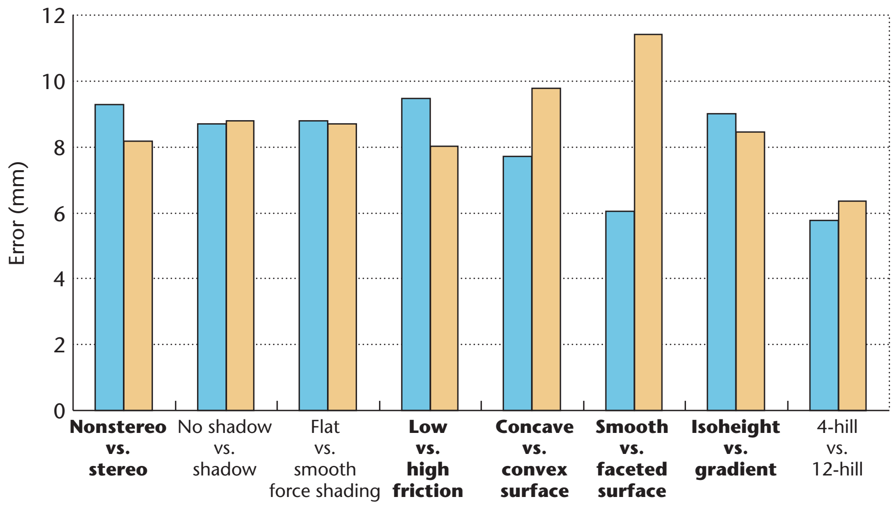

5 Single-factor effects on average-error measure. The statistically significant effects are in bold.

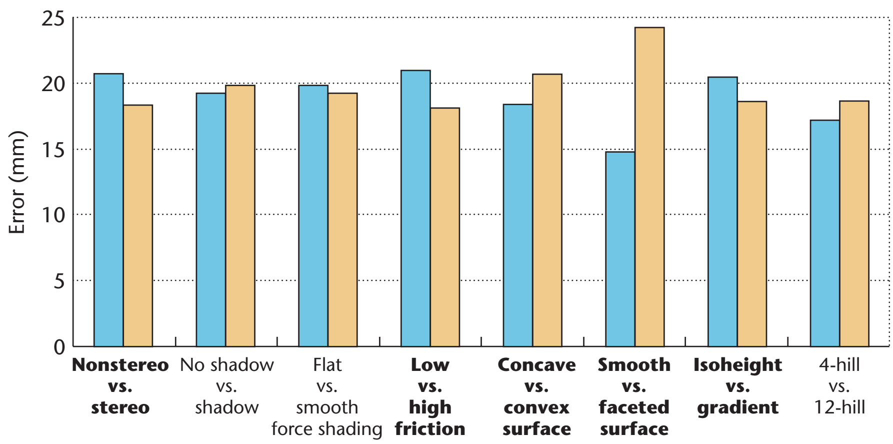

6 Single-factor effects on h_{50} error measure. The statistically significant effects are in bold.

All geometric factors except the hill-frequency factor affected the accuracy considerably. Subjects drew more accurately on concave objects than on convex objects, with the average error reduced by 2.1 mm (21 percent). They also drew more accurately on smooth objects than on objects with discontinuities, with the average error reduced by 5.36 mm (47 percent). Finally, they drew more accurately on target curves that followed the gradient direction of the objects compared to the isoheight curves, with the average error reduced by 0.6 mm (6 percent). The same trends hold for h_{50} worst-case error (see Figure 6).The presence of the paintbrush shadow, the normal force-shading method, and the hill frequency did not have a statistically significant effect on either measure of error.

Of the 13 interactions we examined, two interactions had possible statistical significance under both error metrics. The stereo-convexity (S-C) interaction is significant with $p = 0.04$. As seen in the top left of Figure 7 (next page), stereo displays reduce the error caused by the convexity of the objects. Under nonstereo conditions, the average error difference between curves on concave surfaces and convex surfaces was 2.0 mm, but this difference was reduced to 0.2 mm (a 90 percent reduction) under stereo conditions. The friction-smoothness (T-D) interaction could also be considered statistically significant ($p = 0.09$). As seen in the top right of Figure 7, under the high-friction condition, the disadvantages of nonsmooth objects are reduced. The average error difference of 5.8 mm under the low-friction condition decreased to 4.4 mm (a 24 percent reduction) under the high-friction condition.

Speed and fatigue

The average time for a trial was 23.7 seconds. Stereo, interpolated-normal force shading, smoothness, and hill frequency had statistically significant effects on the speed using a $p < 0.05$ criteria. Figure 8 (on p. 47) shows the single-factor effects on speed. The stereo display improved speed, reducing the average trial time by 1.9 seconds (9 percent) compared to the nonstereo condition ($p = 0.0405$). Although interpolated-normal force shading did not have a statistically significant effect on accuracy, it did have an effect on speed, reducing the average trial time by 3.0 seconds (15 percent) compared to flat shading ($p = 0.011$).

The most statistically significant effects on speed, as well as those with the largest magnitudes, resulted from two geometric factors. Both smoothness and lower hill frequency improved speeds, with $p < 0.0001$. On average, the subjects spent 5.2 seconds less tracing curves on smooth objects than on coarsely faceted versions of

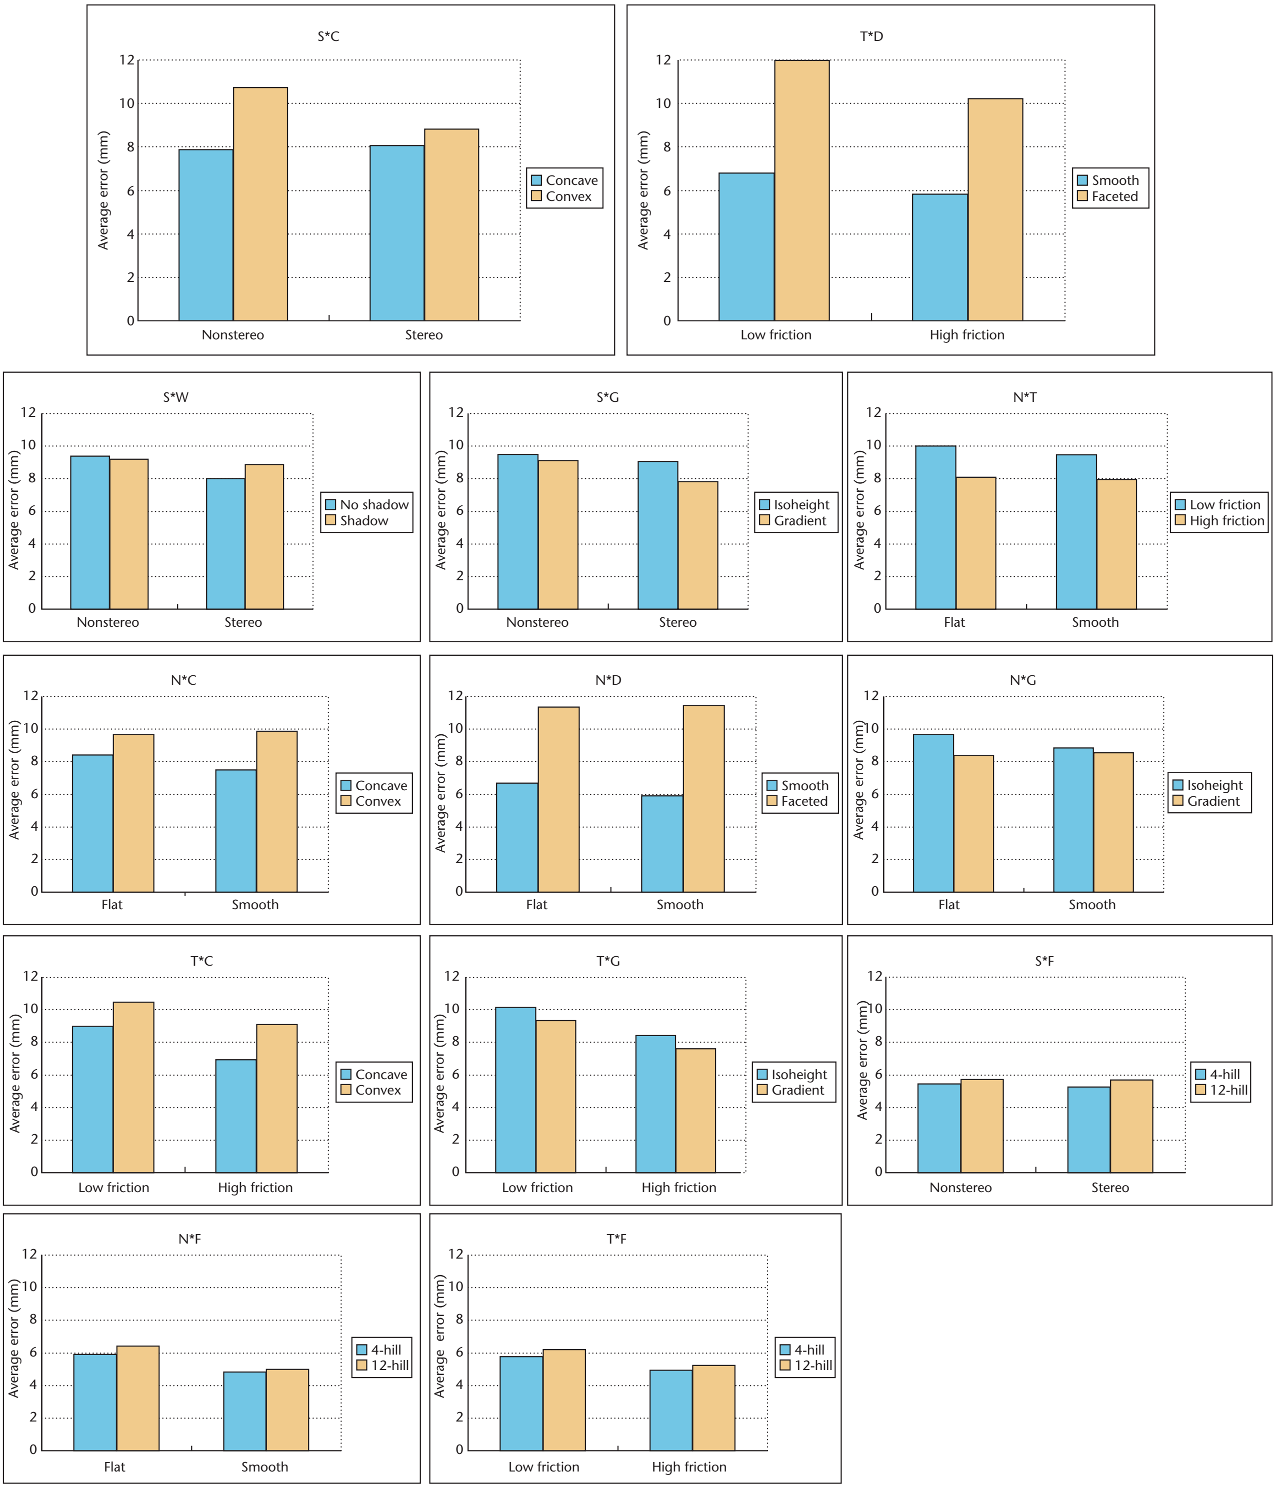

7 **Two factor interactions on average error measure. Only S-C and T-D are statistically significant.**

the same objects, a time reduction of 22 percent. Drawing on the 4-hill surface took 11.8 seconds less on average than on the 12-hill surface, a 34 percent time reduction.

Shadows (W), friction force (T), surface convexity (C), and curve direction (G) did not have a statistically significant effect on speed using a $p < 0.05$ criteria. However, if the criteria were relaxed to $p < 0.1$, the effect of curve direction would be considered statistically significant and the effect of convexity would be borderline. As shown in Figure 9, we found no statistically significant two-factor interactions on the time measure, even using the looser $p < 0.1$ criteria.

To determine if subject fatigue might have affected the results, we compared average error and speed between the two replications of the treatment settings

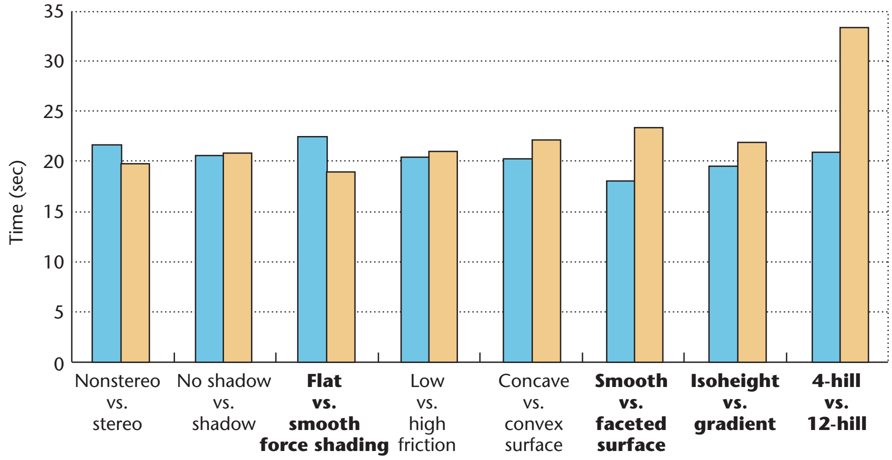

8 Single-factor effects on time measure. The statistically significant effects are in bold.

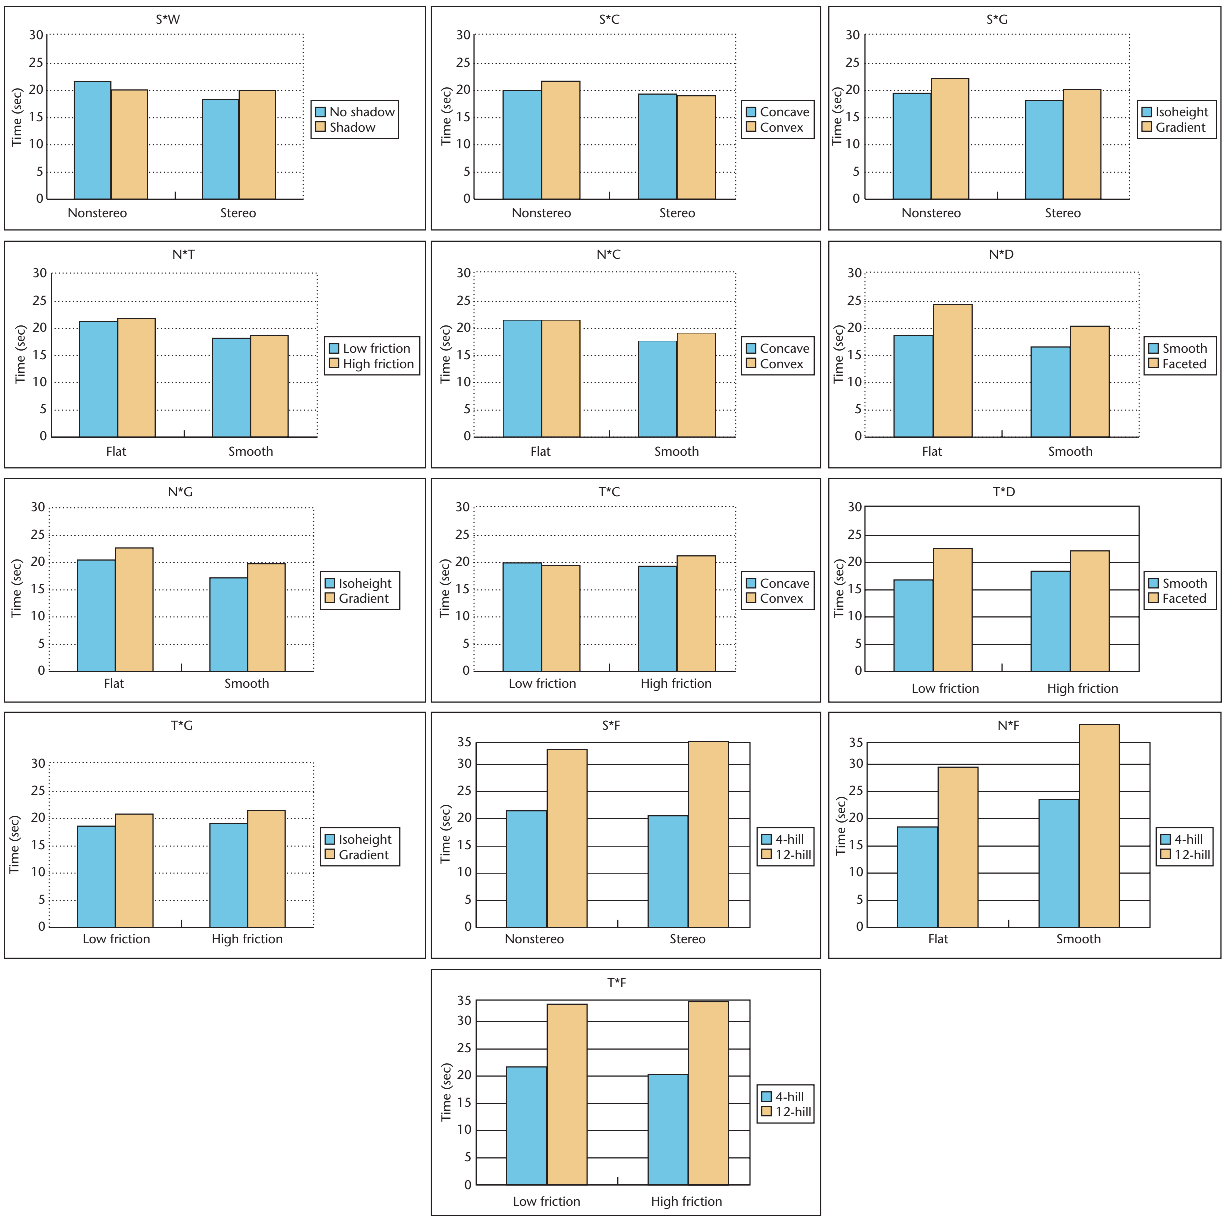

9 Two-factor interactions on time measure. None are statistically significant.

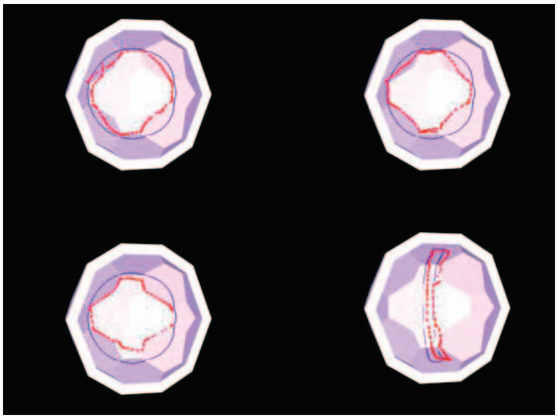

10 **User lines tend to go along the sharp edges on concave surfaces. The blue lines are target curves and the red dots are captured at a constant sampling rate from a user's brush motion.**

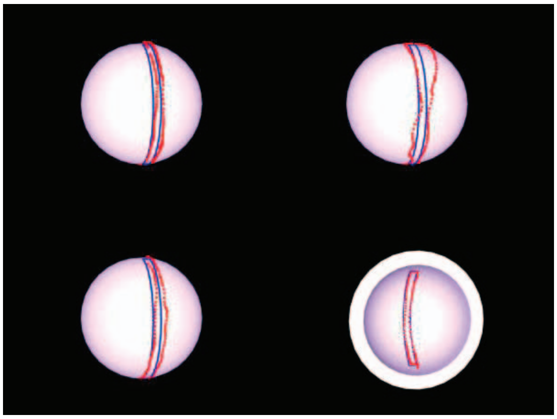

11 **Sparse distribution of sampled points in regions directly facing the users shows that the speed of drawing is faster there than in other regions.**

presented to each user. We tested for statistically significant bias in the differences of 440 pairs of first and second trials using *t*-tests, which test for the statistical significance of the distance between two means. For these tests we removed two outliers outside of the 6σ interval, one a huge improvement and one a huge decline, –45 mm and +70 mm, respectively, compared to $\sigma = 4.2$ mm.

There was a statistically insignificant 0.39 mm increase in average error in the second trial compared to the first trial but a statistically significant average 5.6-second decrease in time for the second trial. The time decrease without the error increase or decrease indicates that fatigue was not negatively affecting user performance, but that we were unable to eliminate the learning effect with the practice sessions. The fact that the order of treatment settings was randomized across subjects made it possible to draw statistically significant conclusions despite this remaining learning effect.

Speed, accuracy, and geometry

We expected that there might be a tradeoff between speed and accuracy for some of the factors we studied. But we did not find an overall correlation, negative or otherwise, and we did not find that a statistically significant effect of a particular factor on speed tended to be accompanied by a statistically significant opposite effect of the same factor on accuracy.

All the geometric factors we tested—including smoothness, convexity, curve direction, and hill frequency—were significant in terms of accuracy or speed. Smoothness was beneficial with strong significance for both. The effects of convexity and curve direction were significant only for accuracy using a criterion of $p < 0.05$, but their probabilities for the null hypothesis for speed were still relatively low ($p = 0.1008$ and 0.0800, respectively). Concave surfaces tended to improve both accuracy and speed. Drawing on the gradient, while more accurate, took longer, the only tradeoff we found between accuracy and speed that was potentially statistically significant. Hill frequency was only statistically significant for speed.

The geometric effects on accuracy tended to have simple, intuitive explanations based on our observations of user behavior. Users' hands tended to slip and move jerkily when they crossed over sharp edges, reducing accuracy for nonsmooth objects. When the sharp edges were on concave surfaces, the lines the subject drew tended to go along the grooves (see Figure 10). Subjects drew more accurately on concave objects, which can stop unwanted motions more easily than convex objects. For instance, there is more likely to be a larger error when the paintbrush crosses a convex edge than when it crosses a concave edge because the concavity tends to stop the brush motion. If these explanations are correct, we would also expect to see an interaction between smoothness and convexity, but unfortunately this is not an interaction that can be reliably measured with our fractional factorial design.

The reason that subjects made larger errors tracing curves whose dominant direction is along the isoheight direction of the objects that they lie on can be explained by the fact that the brush is more likely to slide along the gradient direction accidentally rather than the isoheight direction. Therefore, if the target curve direction is located along the gradient direction, accidental movements are likely to be along the target curve itself. When the target curve is along the isoheight direction, accidental movements are also more likely to be along the gradient direction, but this will be normal to the target curve. The unwanted motions were mostly in the $-z$ direction because the subjects had often been countering force feedback in the $+z$ direction from pushing against the 3D objects before they slipped.

Our data also indicate that the local drawing speed appears to be faster in regions where surfaces face directly toward the subjects. The relatively sparse distributions of points captured at constant time intervals shown in Figure 11 indicates that the drawing speeds are higher near the regions that are facing directly toward the subjects. This fact suggests that there are certain orientations of surfaces that make drawing faster. We plan to study and quantify these speed changes in future research.

The effect of hill frequency has more than one plausible explanation. The more frequent changes in direction when tracing along the 12-hill surface could explain why it took longer than the 4-hill surface, although the higher curvature could also have contributed.

Stereo display and shadows

We found that the stereo display had beneficial effects reducing error and increasing speed. This result is consistent with other research on stereo displays, but it is not consistent with the P.J. Passmore study of stereo with haptic drawing, which found that stereo displays didn't help in the presence of haptics.[8] The interaction we found between stereo and convexity helps explain this apparent discrepancy. In the Passmore study, the curves were all drawn on flat surfaces. Under haptics, the benefits of stereo seem to vary with the geometry.

But it's possible that this effect was not due to the stereo displays alone. Because the effect of the virtual cursor and physical cursor comes only with stereo, this stereo factor (S) should be regarded as stereo-plus pseudo-collocation; collocation may improve performance.[13]

The shadow of the paintbrush was the only factor that did not significantly affect either accuracy or speed. One possible interpretation of this data is that in the presence of a stronger contact cue, such as haptic interfaces, shadows do not provide any added benefit. Another possibility is that shadows are beneficial under some geometric conditions but harmful under other geometric conditions. Our experimental design does not provide for analysis of the interactions between shadows and geometric factors. An alternate explanation is that shadows might be beneficial even in the presence of a haptic interface but only for guiding the brush to an initial contact point on the surface. Because such targeting takes up only a small fraction of our total task time, our experiment was probably not sensitive enough to pick it up.

The high-friction condition makes drawing more accurate without making drawing much slower. High friction seems to help by reducing accidental slips. This interpretation is consistent with the interaction between the smoothness of the objects and the strength of friction. The higher friction seems to eliminate the disadvantage of objects that have sharp edges by reducing accidental paintbrush motion.

The smooth, interpolated-normal force shading improved speed without a significant effect on the error, which again is consistent with our results that interpolated-normal force shading makes nonsmooth objects feel smoother.

Conclusions

This research project verifies that the effects of geometry—including seemingly irrelevant geometric factors such as convexity and gradient direction—are significant for haptic drawing. Thus we emphasize that it's extremely important to study a range of geometric factors when doing user studies for testing haptic interfaces.

We found that by using a high surface-friction force for haptic rendering, disadvantages caused by certain geometry can be overcome. We can prevent jerky and sudden brush motions that can accidentally occur when passing over sharp edges by exerting a high friction force, without decreasing speed by a statistically significant amount. By simulating higher friction forces on our model, we can make surface drawing more accurate.

We also found that stereo displays, in combination with haptics, can provide additional benefits for drawing in terms of reducing error and increasing speed for the drawing task. We found the error reduction with stereo to be most pronounced when drawing on convex surfaces, which had far higher errors without stereo. This suggests that for tasks in VR environments that require accuracy or speed, stereo displays are worthwhile despite their high computing costs, effectively halving the effective frame rate.

Finally, we found that shadows provide negligible benefit in the presence of haptic interfaces. Force feedback appears to be the primary factor in the ability of users to recognize contact with a virtual object. While other studies do show that shadows are helpful for nondrawing tasks, shadow calculations affect frame rate, so the decision to introduce shadows in a VR system should be considered in the context of the expected task mix. ■

Acknowledgments

Portions of this article were taken from a thesis submitted in partial satisfaction of the requirements for an MS at the University of California, Berkeley. We gratefully acknowledge the statistics advice of Lilly Wang and Jing Yi; Mike Montero and other IML researchers for helpful discussions and comments; Carlo Séquin and Paul Wright for support and encouragement; the numerous students who have worked on the CHaMUE system, particularly Cheston Tan for help with the stereo code; Eric Karasuda and the anonymous reviewers for suggestions that improved the exposition; the Stanford 3D Scanning Repository for the bunny model in Figure 2; and all the subjects who participated in the experiment. This work was supported in part by the Ford Motor Company, UC MICRO, and a Junior Faculty Research Grant.

References

1. A. Gaggioli and R. Breining, "Perception and Cognition in Immersive Virtual Reality," *Communication through Virtual Technology: Identity Community and Technology in the Internet Age*, IOS Press, 2001, pp. 71-86.
2. A.D. Gregory, S.A. Ehmann, and M.C. Lin, "inTouch: Interactive Multiresolution Modeling and 3D Painting with a Haptic Interface," *Proc. Virtual Reality Ann. Int'l Symp.*, IEEE CS Press, 2000, pp. 45-52.
3. M. Agrawala, A. Beers, and M. Levoy, "3D Painting on Scanned Surfaces," *Proc. Symp. Interactive 3D Graphics*, ACM Press, 1995, pp. 145-150.
4. Y. Wang and C. MacKenzie, "The Role of Contextual Haptic and Visual Constraints on Object Manipulation in Virtual Environments," *Proc. Conf. Human Factors in Computing Systems*, ACM Press, 2000, pp. 532-539.
5. D. Feygin, M. Keehner, and R. Tendick, "Haptic Guidance: Experimental Evaluation of a Haptic Training Method for a Perceptual Motor Skill," *Proc. 10th Symp. Haptic Interfaces for Virtual Environment and Teleoperator Systems*, IEEE CS Press, 2002, pp. 40-47.
6. L. Wanger, "The Effect of Shadow Quality on the Perception of Spatial Relationships in Computer Generated Imagery," *Proc. Symp. Interactive 3D Graphics*, ACM Press, 1992, pp. 39-42.

7. H. Hu et al., “Visual Cues for Imminent Object Contact in Realistic Virtual Environments,” *Proc. IEEE Visualization Conf.*, IEEE CS Press, 2000, pp. 179-185.
8. P.J. Passmore et al., “Effects of Viewing and Orientation on Path Following in a Medical Teleoperation Environment,” *Proc. Conf. Virtual Reality Ann. Int’l Symp.*, IEEE CS Press, 2001, pp. 209-215.
9. C. Everitt and M. Kilgard, *Practical and Robust Stenciled Shadow Volumes for Hardware-Accelerated Rendering*, tech. report, Nvidia, 12 March 2002.
10. H.B. Morgenbesser and M.A. Srinivasan, “Force Shading for Haptic Shape Perception,” *Proc. ASME Dynamic Systems and Control Division*, ASME, 1996, pp. 407-412.
11. J. Sethian, “A Fast Marching Method” *SIAM Review*, vol. 41 no.2 , 1999, pp. 199-235.
12. G. Taguchi, *Taguchi Methods: A Hands-On Approach*, Addison-Wesley, 1993.
13. C. Ware and J. Rose, “Rotating Virtual Objects with Real Handles,” *ACM Trans. Computer-Human Interaction*, ACM Press, 1999, pp. 162-180.

Readers may contact Youngung Shon at U.C. Berkeley, Dept. of Mechanical Engineering, 2114 Etcheverry Hall MC 1740, Univ. of California, Berkeley, CA 94720-1740; yshon@me.berkeley.edu.

Youngung Shon *is a PhD candidate in mechanical engineering at the University of California at Berkeley. His research interests include haptics, computer graphics, and VR. Shon has an MS in mechanical engineering from U.C. Berkeley.*

Sara McMains *is an assistant professor of mechanical engineering at the University of California at Berkeley. Her research interests include geometric and solid modeling, CAD/CAM, layered manufacturing, computer graphics and visualization, virtual prototyping, and VR. She has a PhD in computer science from U.C. Berkeley. She is a member of ACM Siggraph and ASME.*

For more information on this or any other computing topic, visit our Digital Library at http://www.computer.org/pubs/dlib.