Perceptual Issues in Haptic Digital Watermarking

Domenico Prattichizzo, Mauro Barni, Gloria Menegaz, and Alessandro Formaglio
University of Siena

Hong Z. Tan and Seungmoon Choi
Purdue University

The growing interest in haptic applications suggests that haptic digital media will soon become widely available, and the need will arise to protect digital haptic data from misuse. In this article, we present our study and findings on psychophysical experiments regarding human abilities to perceive a digital watermark, or hidden signal, through a haptic interface.

Haptic interfaces allow physical interactions with virtual 3D objects through the sense of touch. Possible applications include training for minimally invasive or microscopic surgical procedures, interacting with sculptures such as Michelangelo's *David* that we can't directly touch, perceptualizing multidimensional data sets such as earthquake simulations that we can't easily comprehend through visual displays alone, and assistance to sensory-impaired individuals by displaying visual and/or audio information through the haptic sensory channel.

Because of the expected growing importance that digital haptic data will have in the near future, it's easy to predict that the need will soon arise to protect such data from misuse, like unauthorized copying and distribution, or false ownership claims. Among the available technologies to protect digital data, digital watermarking is receiving increased attention, thanks to its unique capability of persistently hiding a piece of information within the to-be-protected data.[1] We can use this hidden information to prove ownership, deny permission of copying the data, or detect tampering.

In this article, we present the results of two psychophysical experiments that investigated the perceptibility and detectability of a hidden signal in the macro- and microgeometry of the virtual object's surface. In the first experiment, we embedded the watermark into a virtual surface's macrogeometry by modifying the underlying 3D model's wireframe. To begin with, we chose a flat surface so that signals related to the object's shape wouldn't inadvertently mask the watermark's detection. Nevertheless, we represented the surface with a 3D mesh so that we could readily extend this initial work to objects with arbitrary surface shapes. We modeled the watermark as an additive white noise superimposed on the host surface. The goal of the experiment was to estimate the noise intensity threshold as a function of the underlying mesh's resolution.

Our second experiment focused on the microgeometry of object surfaces by embedding the watermark in the texture data. We used a simple one-dimensional sinusoidal model for both the watermark and the host signal. The goal of this experiment was to investigate whether existing detection threshold data[2,3] could successfully predict the perceptibility of the watermark. Despite the texture model's simplicity, this experiment provided the first evidence of the possibility of embedding a haptically imperceptible watermark that can later be detected via spectral analysis.

Before going further into these experiments, we first provide a little background information regarding the basic issues and requirements for successful 3D watermarking techniques. (For further information, also see the "Overview of Watermarking Techniques" sidebar.)

Basic issues and requirements

In the past, a great deal of research has focused on digitally watermarking audio, images, and video—meanwhile, haptic interfaces are inherently related to 3D surfaces. Despite the fact that 3D models are widely used in several applications such as virtual prototyping, cultural heritage, and entertainment, watermarking 3D objects is still in its infancy. One of the reasons for this gap lies in the difficulty of extending common signal processing algorithms to 3D data.

The first requirement that any watermarking technique must satisfy is watermark imperceptibility. In the case of still images and video sequences, the imperceptibility requirement has triggered a great deal of research about human visual systems, resulting in a number of possible algorithms that exploit the properties of human vision to improve watermark invisibility while keeping the watermark energy constant.[4]

Recently, we've also seen some progress in 3D watermarking.[5] In this case, the watermark is hosted by the macrogeometry of the considered

1070-986X/07/$25.00 © 2007 IEEE

Published by the IEEE Computer Society

Overview of Watermarking Techniques

Generally speaking, we can consider any watermarking system a communication system consisting of two major components: a watermark embedder and a watermark detector. The watermark usually consists of a pseudorandom sequence with uniform, binary, or Gaussian distribution. It's transmitted through the watermark embedder over the original, to-be-marked object (in our case, a 3D surface). The watermark detector extracts the watermark from the marked data. Intentional and unintentional attacks and distortions applied to the mesh hosting the watermark further characterize and complicate the transmission channel.

We can divide watermarking techniques into two categories:

- spatial/temporal domain techniques that directly add the watermark to pixel values and
- transformed domain techniques that add the watermark in the frequency domain.

Once we've chosen the host features, we need to specify the embedding rule. The most common approach to watermark embedding is the additive rule according to which $y_i = x_i + \gamma w_i$, where x_i is the ith component of the original feature vector, w_i the ith sample of the watermark, γ a parameter controlling the watermark strength, and y_i the ith component of the watermarked feature vector.

Recently, a new approach to watermark embedding has been proposed. This approach, commonly referred to as informed watermarking or Quantization Index Modulation (QIM) watermarking,[1] can greatly improve the system's performance as a whole. However, for the sake of simplicity, our analysis focused on additive watermarking, leaving the analysis of QIM schemes for future work.

The way the watermark is extracted from data plays a crucial role. In blind decoding, the decoder doesn't need the original data (mesh) or any information derived from it to recover the watermark. Conversely, nonblind decoding refers to a situation where extraction is accomplished with the aid of the original, nonmarked data. We can also make an important distinction between algorithms embedding a mark that can be read and those inserting a code that can only be detected. In the former case, we can read the bits contained in the watermark without knowing them in advance. In the latter case, we can only verify the bits if a given code is present in the document. Though our perceptibility analysis was very general, we specifically focused on the case of blind watermark detection.

As mentioned in the introduction, an important aspect of any watermarking system is the imperceptibility of the hidden information. For this reason it's of primary importance that we carefully study the properties of the sensory modality through which subjects detect the marked data.

In audio watermarking, researchers have exploited existing data from studies on the human auditory system to better hide the watermarking signal within the host audio. More relevant to the 3D scenario is the case of still image watermarking. Several models of the human visual system have been modified and exploited to ensure the hidden signal's invisibility.

In most cases, Watson's simple model of vision has been adopted,[2] leading to watermarking systems working in the discrete cosine transform (DCT) or discrete Fourier transform (DFT) domain. Watson's model is able to predict the visibility of a sinusoidal grating (watermarking signal) superimposed on another sinusoidal grating (host signal).

One problem with visual watermarking in the frequency domain is the lack of spatial localization—hence, alternative models operating in the wavelet domain have been proposed that have led to improved watermark invisibility. As far as 3D meshes are concerned, few studies have been published so far. Of those studies, two different approaches have been taken: judging the watermark's visibility in selected views of the rendered mesh, and allowing the observer to freely play with the mesh by zooming and rotation.[3]

Researchers still have much work to do to fully understand and successfully implement watermark visibility in 3D objects. To the best of our knowledge, no previous work on haptic watermark perceptibility has been presented, with the exception of the studies carried out by ourselves.

For a more detailed discussion of watermarking issues, we refer readers elsewhere.[2,4]

References

1. B. Chen and G. Wornell, "Quantization Index Modulation: A Class of Provably Good Methods for Digital Watermarking and Information Embedding," *IEEE Trans. Information Theory*, vol. 47, no. 4, 2001, pp. 1423-1443.
2. I.J. Cox, M.L. Miller, and J.A. Bloom, *Digital Watermarking,* Morgan Kaufmann, 2001.
3. M. Corsini et al., "Objective Evaluation of the Perceptual Quality of 3D Watermarking," *Proc. Int'l Conf. Image Processing,* IEEE CS Press, 2005, pp. 241-244.
4. M. Barni and F. Bartolini, *Watermarking Systems Engineering: Enabling Digital Assets Security and Other Applications*, Marcel Dekker, 2004.

virtual object's surface, which we can assume is represented by a 3D mesh. Accordingly, we can judge the watermark's intrusiveness in terms of its visibility in the mesh's rendered version. More generally, in applications where the virtual object is sensed through a haptic interface, guaranteeing the imperceptibility of the watermark requires characterizing the haptic channel's sensitivity.

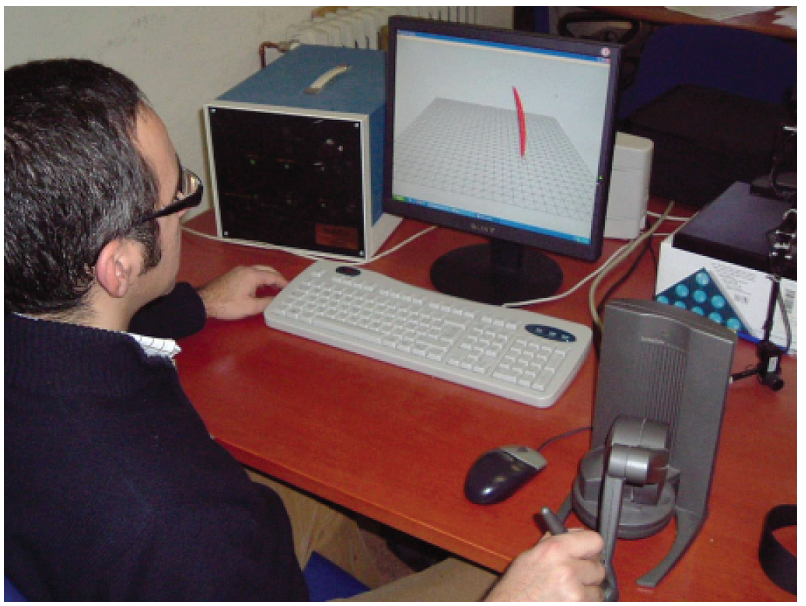

Figure 1. Subject touching a virtual surface through a stylus-like device.

Despite an exponential increase in haptics research activities in the last decade, our understanding of how people sense and manipulate objects with their hands is still limited.[6] The most popular haptic interfaces—such as SensAble Technologies' Phantom and Force Dimension's Delta—let us interact with the virtual environment through one contact point only. Interfaces with a higher number of degrees of freedom (DOF) and with multiple interaction points are available, but are less common or reliable than those with three DOF and one interaction point.

With the term *haptic rendering*, we refer to a branch of haptics research that deals with the calculation of interaction forces between a virtual representation of the user and a virtual object. In most cases, haptic rendering is a two-step process consisting of shape- and texture-based force rendering. In this context, shape refers to the macrogeometry of an object's surface, as opposed to texture that describes the surface's fine structure, or microgeometry. To render an object's shape, we can use typical single-point contact rendering algorithms such as the god–object algorithm.[7] To render the texture of a virtual surface, we can perturb the shape-based force using a texture model.[8]

Experiment 1: Macrogeometry watermarking

With our first experiment, we hoped to estimate the perceptibility threshold of a watermark modeled as white noise with uniform distribution embedded in the surface's macrogeometry description. As we previously noted, we began by using the simplest case of a flat surface implicitly described by a 3D mesh. We used the same representation for both the host plane and the watermark. The 3D meshes were encoded in data structures representing the spatial coordinates of all the vertices as well as their interconnections. We haptically displayed a virtual mesh using a force-feedback device that allowed single-point contact mediated by a stylus, as Figure 1 depicts. We conveyed information about the surface shape via the direction of the reaction forces that corresponded to the normal vectors to the mesh. The force interaction model did not include friction.

We embedded the digital watermark in the surface's macrogeometry by modifying the data matrices according to the additive rule. In this case, we added the watermark signal to the height of the corresponding mesh's vertex. The watermark's strength was represented by the noise spectral power of the equivalent noise model. Human sensitivity to the noise was estimated as the minimum noise level required for the watermark to become detectable. Since the mesh's resolution—that is, the dimension of triangle elements—could vary with application specifications and surface shapes, we conducted the experiment using several 3D meshes with different resolutions. In this way, we were able to establish the relationship between the sensitivity to the watermark's strength and the size of the triangular mesh elements.

Methods

The host surface was a horizontal square plane of size 15×15 cm^2 represented by a 3D triangular mesh and placed in front of the subject. Let $\mathbf{v}(i)$ be the 3D vector of the ith triangle vertex and $\mathbf{n}(i)$ the surface normal defined at this point. As we mentioned earlier, the embedded watermark altered the mesh vertices according to the following rule:

$$\mathbf{v}_w(i) = \mathbf{v}(i) + w(i)\mathbf{n}(i),$$

where $\mathbf{v}_w(i)$ was the ith watermarked vertex and $w(i)$ the watermark noise model. Specifically, we assumed a uniform distribution for $w(i)$ in the range $\{-\Delta, +\Delta\}$. The corresponding frequency domain representation of the watermark noise consisted of a constant spectral power over all frequencies, $P_w(\omega) = \Delta^2/12$.

We asked five human subjects to explore the virtual surfaces using a desktop model of the Phantom force-feedback device. They held the Phantom's stylus with their right hand and

stroked the surface, as Figure 2 illustrates.

We used an impedance model[7] to render a force F to the subject's hand (see Figure 2) when the stylus tip was inside the virtual surface. No force was displayed when the stylus was outside the virtual surface.

We used a two-interval forced choice (2IFC) three-down, one-up adaptive procedure to estimate the watermark detection threshold.[9] According to the 2IFC paradigm, there were two stimulus alternatives, one with the host mesh, and the other with the watermarked mesh. In each trial, we presented the two surfaces to the subject in random order. The subject's task was to report which surface (the first or the second) contained the plane with the watermark. As is typical of most adaptive procedures, we did not provide trial-by-trial correct-answer feedback during the experiment. According to the three-down, one-up adaptive rule, the watermark strength was decreased after three consecutive correct answers and increased after a single wrong answer, as follows:

$$\begin{cases} \Delta(0) = 2\ mm \\ \Delta(i+1) = 0.5\Delta(i) & \text{(after 3 correct responses)} \\ \Delta(i+i) = 1.5\Delta(i) & \text{(after 1 incorrect response)} \end{cases}$$

where the initial value $\Delta(0) = 2$ mm was found to be clearly perceivable in a pilot test.

We terminated each experimental run after six reversals. A reversal occurred when the watermark strength changed from increasing to decreasing or vice versa. It's worth noting that the total number of trials per run was not fixed a priori, but was determined adaptively to meet the stop condition. Figure 3 shows a sketch of a typical staircase sequence produced during one experimental run. The detection threshold was computed by taking the average of the peaks and valleys over the six reversals within one staircase sequence.

The experiments were arranged in two blocks per subject: a practice block and an experimental block. Each block consisted of seven runs corresponding to the seven sizes of the side of the triangular mesh elements ranging from 2 to 10 mm. We recorded the estimated detection thresholds from the second block for each subject as a function of mesh resolution. Five subjects, aged between 22 and 25, participated in the experiment. All were right-handed with no known sensorimotor impairments. Their prior experience with the Phantom device varied from naive to expert.

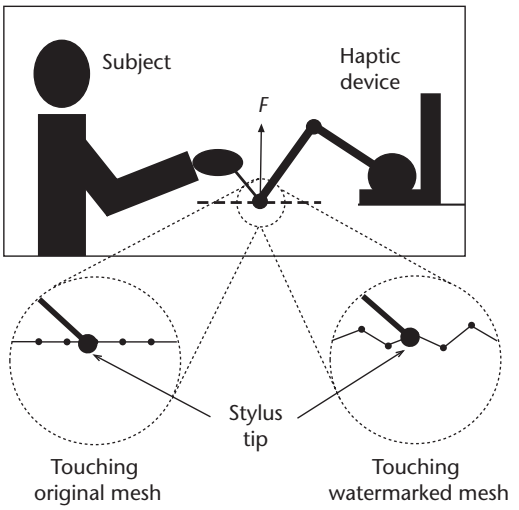

Figure 2. Experimental setup. The bottom left is a host mesh and the bottom right is a watermarked mesh.

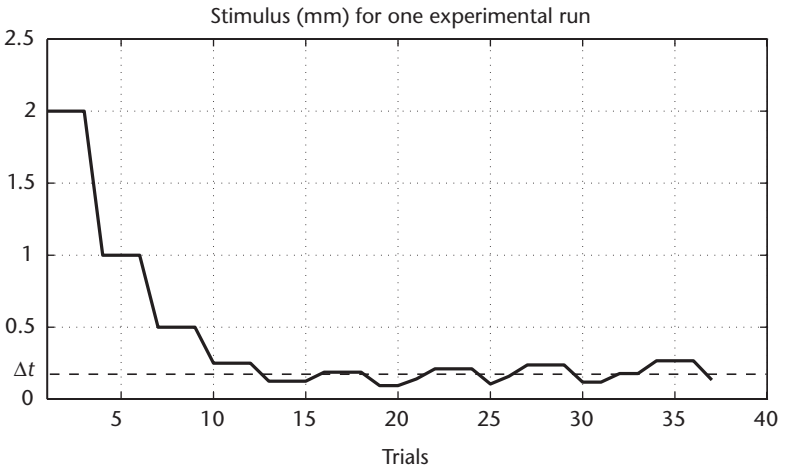

Figure 3. A typical staircase sequence for one experimental run. In this case, the triangle side length was $l = 2.85$ mm. We determined the watermark strength by the parameter Δ. The dashed line represents the estimated detection threshold.

Results and discussion

We calculated the mean and standard deviation of the estimated detection thresholds in terms of the watermarking strength Δ from the data of all the subjects. Figure 4 (next page) shows the average thresholds as a function of the triangle side length l. Note that the procedure we followed was a within-subject design, meaning that we tested each subject with all values of the triangle mesh side length. As a result, we were able to address the effect of side length on the watermark detection threshold with each subject. Because of the relatively small number of subjects tested and their different levels of prior experience with haptic interfaces, it's fair to say that there could be some between-subject variability in thresholds, thereby explaining the relatively large standard deviations in the plot of Figure 4.

In general, watermark detection threshold increased as a function of the mesh resolution,

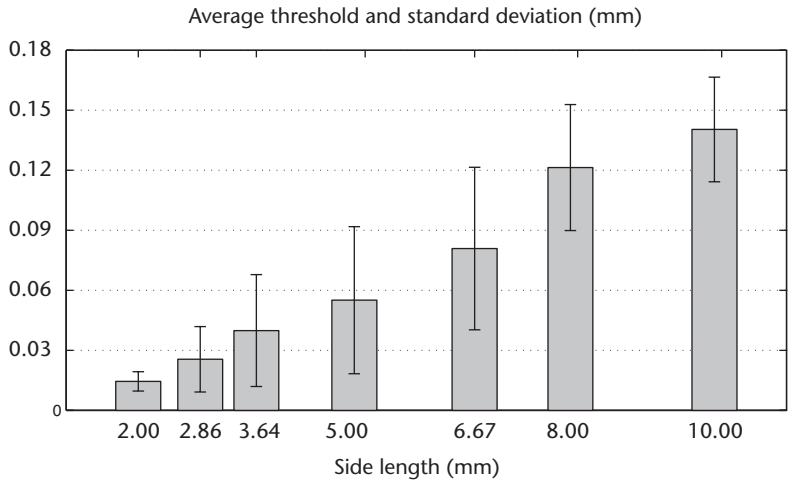

Figure 4. Average thresholds over triangle side length.

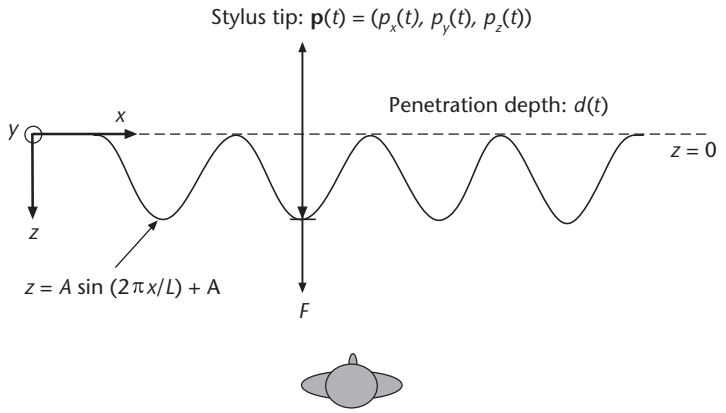

Figure 5. Bird's-eye view of subject, textured vertical plane, and coordinate frame. The dashed line indicates the flat vertical plane upon which we superimposed a one-dimensional sinusoidal texture model. Subjects stroked the textured surface along the x-axis. We measured the penetration depth as the distance between the stylus tip and the point on the textured surface along the z-axis.

indicating that we can embed more noise in surfaces with coarser representation. We can also use the values of the thresholds to adjust the watermark signal's strength as a function of the local geometrical features of the host surface, so that we can guarantee imperceptibility.

Experiment 2: Microgeometry watermarking

Our second experiment considered the watermarking of haptic virtual textures—that is, the microgeometry of surfaces.[10] While the first experiment aimed to estimate a human's sensitivity to watermarks as a function of the size of the macrogeometry's triangular mesh, the second experiment's goal was to investigate the exploitability of existing data in the literature for predicting the perceptibility of watermarks embedded in the microgeometry.

Past research on human detection thresholds for sinusoidal stimuli[2,3] has established the minimum signal strength required for producing a sensation. Because we chose to model haptic virtual textures (both the host and the watermark signals) using sinusoidal gratings, we anticipated that we could predict the perceptibility of watermarks using existing detection thresholds.

When we explore a virtual flat haptic surface with a superimposed sinusoidal grating (texture) with a force-feedback device such as the Phantom, the device conveys texture information through vibration. Previous work has shown that the temporal signal contributing to texture perception is characterized by a spectral peak of the force or position signals recorded near or at the stylus tip.[11] We can determine the frequency of this peak by the spatial period of the sinusoidal grating and the speed at which the textured surface is stroked. The peak's amplitude determines the texture's perceived intensity (or roughness).

It therefore follows that we can consider the digital watermark for a virtual haptic texture in the spectral domain. Given a host's texture signal, we can add a spectral peak at a different frequency, with an amplitude below the human detection threshold (the watermark) that guarantees its imperceptibility. Therefore, the second experiment employed a simplified version of the additive watermarking method outlined earlier.

Methods

We defined the height map of the host texture signal by

$$h(x) = A_h \sin\left(\frac{2\pi}{L_h}x\right) + A_h$$

where $A_h = 1$ mm and $L_h = 2$ mm. The symbols A and L denoted the amplitude and the spatial wavelength of the sinusoidal gratings respectively. We defined the watermarked texture signal by

$$h(x) = A_h \sin\left(\frac{2\pi}{L_h}x\right) + A_h + A_w \sin\left(\frac{2\pi}{L_w}x\right) + A_w$$

where L_w = 5 mm, and A_w was either 0.2 (condition 1) or 0.5 mm (condition 2). Figure 5 illustrates the one-dimensional sinusoidal texture model. As in the first experiment, we computed the feedback force F in Figure 5 according to the impedance model.[7]

In the spatial domain, the watermarked texture signal was a modulated sinusoidal signal (see the bottom trace in Figure 6). In the frequency domain, it exhibited two spectral peaks (see Figure 7). The upper panel in Figure 7 shows the spectral density (the solid line) of $p_z(t)$ for condition 1, where the weaker watermark signal was embedded in the host texture signal. The $p_z(t)$ data were recorded from a single stroke of the watermarked, textured surface using the Phantom haptic device. The dashed line in the same panel shows the human detection thresholds taken from the literature.[2]

The two spectral peaks corresponded to the watermark (~40 Hz) and host (~76 Hz) signals respectively. The lower panel in Figure 7 shows the same for condition 2, where we used the stronger watermark signal. We were able to predict the perceptibility of the two watermarks by comparing the watermark peaks with the corresponding human detection thresholds. Because the peak for the weaker watermark was at roughly the same level as the human detection threshold, we didn't expect subjects to be able to detect it. The peak for the stronger watermark, however, was clearly above the human detection threshold. We therefore expected that subjects could easily perceive this watermark.

We used a one-interval, two-alternatives, forced-choice paradigm to measure the subject's ability to discriminate the host texture from the watermarked texture. On each trial, the subject felt either the host texture alone, or the host texture with the watermark. Each subject's task was to respond to the host texture and the watermarked host texture.

No trial-by-trial correct-answer feedback was provided during data collection. Each subject performed four 100-trial blocked runs, two for condition 1 and two for condition 2. The order of the four runs was randomized for each subject. At the beginning of each run, subjects familiarized themselves with the stimuli by entering either 1 or 2 on a keyboard to feel the corresponding texture. Training was terminated by the subjects whenever they were ready.

Data from each condition formed a 2×2 stimulus-response matrix consisting of 200 trials.

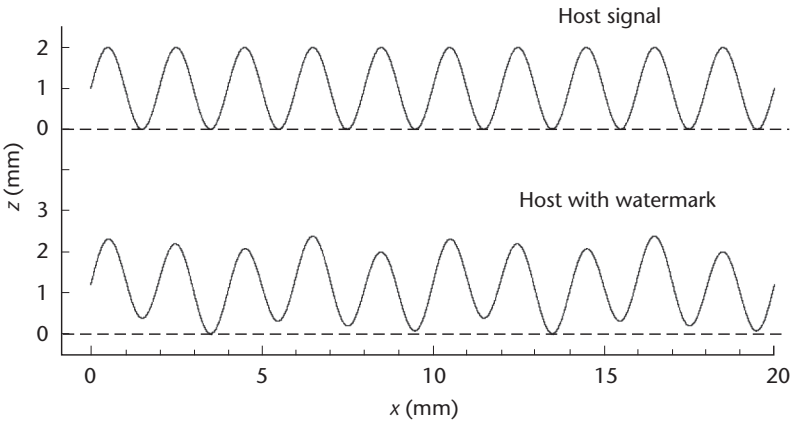

Figure 6. Spatial representation of stimuli. The top trace shows the z versus x sinusoidal grating for the host texture alone (A_h =1 mm, L_h =2 mm). The bottom trace shows the same host signal with an embedded watermark (A_w = 0.2 mm, L_w = 5 mm).

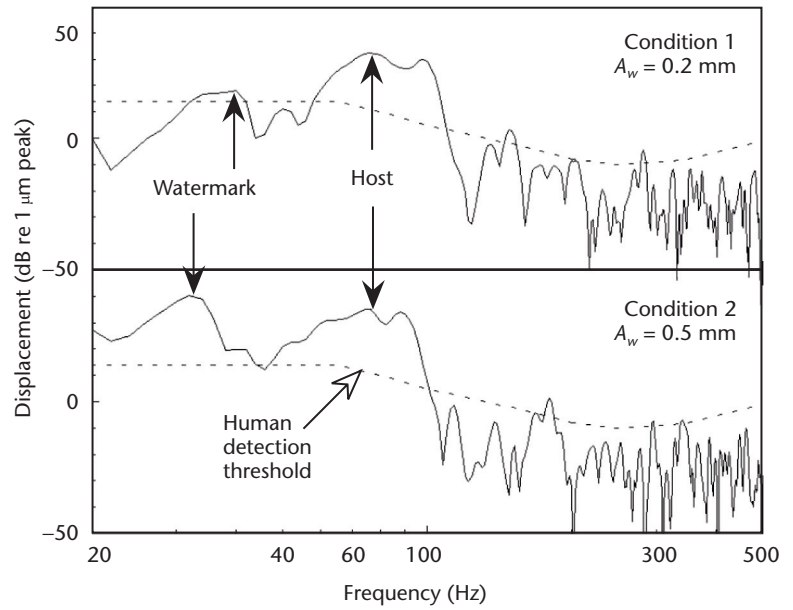

Figure 7. Power spectral densities of $p_z(t)$ for the two watermarked textures (solid lines) and human detection thresholds (dashed lines). The upper and bottom panels correspond to the weaker and stronger watermarks respectively. Arrows indicate the locations of the spectral peaks corresponding to the watermark and the host textures.

Instead of calculating the percent-correct scores that are often confounded by subjects' response biases, we estimated the sensitivity index d' that provided a bias-free measure of the discriminability between the host and watermarked host textures (that is, the perceptibility of the watermark signal).[12,13] A d' value of 0.0, 1.0, or 2.0 cor-

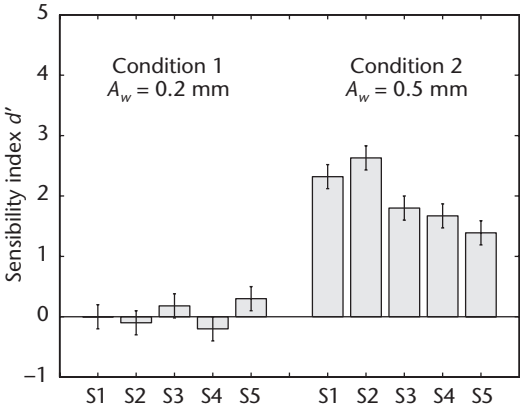

Figure 8. Experimental results. Shown are the sensitivity indices and the corresponding standard deviations for subjects S1 through S5 under the two watermarking conditions.

responds to a percent-correct score of 50, 69, or 84, respectively, assuming no response biases.

Five subjects, aged 25 to 39, participated in the experiment. All were right-handed with no known sensorimotor impairment with their hands. Their prior experience level with the Phantom device varied from naive to expert.

Results and discussion

Figure 8 shows the values of sensitivity index *d'* for five subjects. The *d'* values were essentially 0 in condition 1, where we used the weaker watermark signal—indicating that the subjects couldn't tell the difference between the host texture alone and the watermarked texture. In condition 2, where we used the stronger watermark signal, the values of *d'* were in the range of 1.39 to 2.63—indicating high discriminability. Therefore, the stronger watermark signal was clearly perceivable to all the subjects.

The spectral-domain analysis of the texture signals effectively provided a means for detecting watermarks embedded in a texture signal. The frequency of the host or the watermark texture signal is around $|v|/L$, where v is the average stroking velocity and L the spatial period of the sinusoidal grating.[11] We estimated the average stroking velocity from the position data along the lateral stroking direction—as you can see with $p_x(t)$ in Figure 5. From there, we were able to look for a spectral peak near $|v|/L_w$ to detect the watermark.

Conclusions and future work

We've taken a first step toward analyzing the haptic perceptibility of 3D digital watermarks, both at micro- and macrogeometry levels. As is the case for any novel interdisciplinary research framework, many issues are left open. Many aspects deserve further investigation, including the use of more complex shapes and different models for the watermarking signal, as well as considering the perceptual impact of different rendering techniques.

We also plan to compare haptic and visual perceptibility using the same object and surface representations. This will help us analyze whether the constraints set by the haptic channel are more or less stringent than those set by the visual channel, and if these constraints follow the same rules in both domains.

To avoid reinventing the wheel, we'll systematically test the perceptibility of common visual watermarking techniques in haptically rendered 3D objects at both macro- and microgeometry levels, comparing the perceptibility thresholds for the visual and haptic sensory modalities. To the extent that some of the existing watermarking techniques can be readily applied to the haptics domain and possibly result in higher thresholds (that is, where it's harder to perceive watermarks) via sense of touch, we hope to find new ways to achieve multimodal imperceptibility by employing existing visual watermarking algorithms. It's also quite possible that lower detection thresholds could be found using multimodal (visual and/or haptic) interfaces, in which case we need to develop new watermarking techniques.

Characterizing the two sensory modalities under both unimodal (haptic or visual stimulation alone) and bimodal (visuohaptic stimulation) conditions will help us determine which sensory modality and/or stimulation condition ultimately sets the boundary for the detectability of haptic/visual watermarks. We envision that with the availability of lower-cost, commercially available haptic interfaces, the area of haptic and multimodal digital watermarking will soon become the next fruitful territory for research on digital watermarking. **MM**

Acknowledgments

This work was supported in part by Fondazione Monte dei Paschi di Siena and by the Universitá di Siena under grant PAR2003 and PRIN-2006 AIDA, and by US National Science Foundation awards under grants 0098443-IIS and 0328984CCF.

References

1. M. Barni and F. Bartolini, "Data Hiding for Fighting Piracy," *IEEE Signal Processing*, vol. 21, no. 2, 2004, pp. 28-39.

2. R.T. Verrillo, "Effect of Contactor Area on the Vibrotactile Threshold," *J. Acoustical Soc. of Am.*, vol. 35, 1963, pp. 1962-1966.
3. J.M. Weisenberger and M.J. Krier, "Haptic Perception of Simulated Surface Textures via Vibratory and Force Feedback Displays," *Proc. ASME Dynamic Systems and Control Division*, Am. Soc. Mechanical Eng., 1997, pp. 55-60.
4. R.B. Wolfgang, C.I. Podilchuk, and E.J. Delp, "Perceptual Watermarks for Digital Images and Video," *Proc. IEEE*, vol. 87, no. 7, 1999, pp. 1108-1126.
5. M. Corsini et al., "Objective Evaluation of the Perceptual Quality of 3D Watermarking," *Proc. Int'l Conf. Image Processing,* IEEE CS Press, 2005, pp. 241-244.
6. J.K. Salisbury, "Making Graphics Physically Tangible," *Comm. ACM,* vol. 42, no. 8, 1999, pp. 74-81.
7. C. Zilles and K. Salisbury, "A Constraint-Based God-Object Method for Haptic Display," *Proc. IEEE/Robotics Soc. of Japan Int'l Symp. Intelligent Robots and Systems*, vol. 3, 1995, pp. 146-151.
8. T.H. Massie, "Initial Haptic Explorations with the Phantom: Virtual Touch through Point Interaction," master's thesis, Dept. Mechanical Eng., Massachusetts Inst. of Technology, 1996.
9. H. Levitt, "Transformed up-down Methods in Psychoacoustics," *J. Acoustical Soc. of Am.*, vol. 49, no. 2, 1970, pp. 467-477.
10. D. Prattichizzo et al., "Perceptibility of Haptic Digital Watermarking of Virtual Textures," *Proc. World Haptics Conf.*, IEEE Press, 2005, pp. 50-55.
11. S. Choi and H.Z. Tan, "Perceived Instability of Virtual Haptic Texture. I. Experimental Studies," *Presence: Teleoperators and Virtual Environments*, vol. 13, 2004, pp. 395-415.
12. N.A. Macmillan and C.D. Creelman, *Detection Theory: A User's Guide*, Lawrence Erlbaum Assoc., 2nd ed., 2004.
13. T.D. Wickens, *Elementary Signal Detection Theory*, Oxford Univ. Press, 2002.

Domenico Prattichizzo is an associate professor of robotics and automation at the University of Siena, Italy. His research interests are in the areas of visual serving, robotic grasping, haptics, and geometric control. Prattichizzo received an MSc in electronics engineering and a PhD in robotics and automation from the University of Pisa, Italy.

Mauro Barni is an associate professor at the University of Siena. His research focuses on several aspects of multimedia security, including digital watermarking and digital rights management. Barni received a PhD in informatics and telecommunications engineering from the University of Florence, Italy. He is a senior member of the IEEE.

Gloria Menegaz is an adjunct information engineering professor at the University of Siena. Her research field is perception-based image processing. Menegaz received an MSc in electrical engineering and an MSc in information technology from the Politecnico di Milano, Italy, and a PhD in applied sciences from ITS-EPFL, Switzerland. She is a member of the IEEE, the SPIE, and Eurasip.

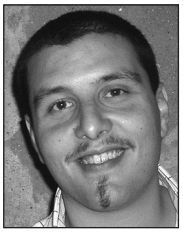

Alessandro Formaglio is a PhD student at the Department of Information Engineering, University of Siena. His research activities deal with haptically enabled virtual reality, especially with haptic perception and virtual object manipulation.

Hong Z. Tan is an associate professor of electrical and computer engineering at Purdue University. Her research focuses on haptic human–machine interfaces and haptic perception. Tan received an SM and a PhD in electrical engineering and computer science from the Massachusetts Institute of Technology. She is a senior member of the IEEE.

Seungmoon Choi is an assistant professor in computer science and engineering at POSTECH, South Korea. His research interests include haptic rendering, psychophysics, and robotics. Choi received a PhD in electrical and computer engineering from Purdue University. This work was performed while Choi was a postdoctoral research associate at Purdue.

Readers may contact Alessandro Formaglio at formaglio@dii.unisi.it.