

Using Multidimensional Scaling to Quantify the Fidelity of Haptic Rendering of Deformable Objects

Peter Leškovský*
Computer Vision Laboratory
ETH Zürich
8092 Zürich, Switzerland

Theresa Cooke†, Marc Ernst
Max Planck Institute for Biological Cybernetics
Tübingen, Germany

Matthias Harders
Computer Vision Laboratory
ETH Zürich
8092 Zürich, Switzerland

ABSTRACT

In this paper, we examine the application of a psychophysical evaluation technique to quantify the fidelity of haptic rendering methods. The technique is based on multidimensional scaling (MDS) analysis of similarity ratings provided by users comparing pairs of haptically-presented objects. Unbeknownst to the participants, both real and virtual deformable objects were presented to them. In addition, virtual objects were either presented under higher-fidelity rendering condition or under lower-fidelity condition in which force filtering and proxy-point filtering were removed. We hypothesized that reducing fidelity of virtual rendering would exaggerate the difference between real and virtual objects. MDS analysis of pairwise similarity data provided quantitative confirmation that users perceived a clear difference between real and virtual objects in the lower-fidelity, but not in the higher-fidelity condition. In the latter, a single perceptual dimension, corresponding to stiffness, sufficed to explain similarity data, while two perceptual dimensions were needed in the former condition. This study demonstrates how MDS analysis provides an opportunity to visualize and quantify the perceptual effects of changes in rendering parameters and how it can be used in the evaluation of haptic rendering scenarios.

Keywords: haptic rendering, haptic Turing test

1 INTRODUCTION

Haptic rendering of deformable objects is a critical element of several multimodal virtual reality applications which are used for instance in medical, entertainment, and education sectors. The main tradeoff in selecting appropriate rendering approaches for deformable objects is that of real-time capability vs. fidelity. This often leads to the key question of how realistic a computer-based approach can be in simulating a real object. One approach to determining the fidelity of a rendering method is to perform a Turing-like test, comparing the haptic feedback which users obtain from real objects against that which they receive from virtual objects approximating real behaviour. Here, we present a specific approach and analysis technique for such a test based on multidimensional scaling (MDS) and demonstrate how MDS can be used to quantify perceptual differences in haptic rendering techniques.

MDS techniques are a family of algorithms which take pair-wise proximities between pairs of objects (e.g., distances between pairs of cities) as input and return the coordinates of the objects embedded in a multidimensional space (e.g., a geographical map showing the cities' relative positions). MDS has been used extensively by psychologists in the study of a wide range of multivariate stim-

uli, allowing for the identification of important psychological dimensions of stimulus variation and quantification of perceptual distances between stimuli, e.g., [24, 19, 15, 16]. Using human similarity ratings on a set of objects as the input proximity data, the output configuration can then be interpreted as a map of the objects in a psychological space which explains the similarity data [1]. MDS provides information about the *dimensionality* of the psychological representation (i.e., how many stimulus dimensions play a role in human similarity ratings and what the dimensions' relative contributions are), as well as more specific topological information about the representations, such as the relative spacing between objects along a given dimension.

In this study, we used MDS to assess the fidelity of haptic interactions. As proximity data, we used similarity ratings provided by users who haptically probed a set of real and virtual objects. Using MDS, we wished to test whether real and virtual objects varying in stiffness would be separable along one or more perceptual dimensions. For a truly high-fidelity interaction, real and virtual objects should be indistinguishable and, if the objects vary solely in terms of stiffness, MDS should recover a single perceptual dimension corresponding to perceived stiffness. For lower-fidelity interactions, more than one perceptual dimension may emerge from the interaction, e.g., a dimension along which real and virtual stimuli are separated.

The paper is organized as follows: after a review of related work, Section 3 introduces the virtual deformation model tested here and describes the selection of deformation parameters based on reference silicone objects. The haptic system used in the experiment is presented in Section 4. In Section 5, the similarity ratings procedure and MDS analysis is described, followed by results and conclusions from the application of this similarity-based approach to validation.

2 RELATED WORK

Measuring the efficiency, accuracy and realism of haptic virtual environments is a challenging task since it requires consideration of device and user behavior, evaluation of the force model computation and of the entire force rendering pipeline, as well as an assessment of the perceptual psychophysics. Despite the need to validate the entire haptic virtual reality system, most studies have evaluated separate components as opposed to carrying out a holistic evaluation.

The suitability of a specific haptic device for use within a haptic scenario can be judged based on a number of physical device characteristics. [14] reported various performance measures which should be provided by the manufacturer and are important for a detailed description of the device's behavior. However, without a user's perceptual validation of the generated force feedback, these measures do not provide sufficient information for an overall description of the device's performance once it is incorporated into a virtual reality system.

Both quantitative performance measures and subjective user preferences were used to evaluate haptic feedback in [12]. Sev-

*e-mail: pleskovs@vision.ee.ethz.ch

†e-mail: theresa.cooke@tuebingen.mpg.de

eral haptic guidance approaches were compared within a steered path-following scenario. The user’s performance was measured by a mean square error of the deviations of the steered object’s locations from a settled course, sampled throughout the interaction time. In the analysis, the authors tested for an effect of interaction factors (i.e., path complexity, visibility and guidance method) on user performance. Additionally, qualitative insight into the different haptic cues provided was obtained by gathering user ratings of the perceived level of control and preferred guidance method. Although this study presented a quantitative performance evaluation and a qualitative assessment of a haptic algorithm, the nature of the generated force feedback was synthetic and thus evaluating the fidelity of the force feedback was irrelevant.

Rating the realism of the haptic feedback within a virtual-reality environment is a secondary aspect for many studies. This can be explained by the nature of the haptic application itself: haptic rendering often involves the generation of artificial forces that are not comparable to natural interactions with real-world objects. Therefore, a high-quality force feedback is not needed. To assess the fidelity of a simulation algorithm, as compared to a real-world phenomenon, standard methods from experimental physics can be used. Depending on what the aspects of a simulated object interaction are, different data (i.e. positions, forces, time, etc.) recorded during experiments carried out in a real-world setting can be set against data gathered in virtual simulations. An example of such a comparative assessment was presented in [10]. Within their study, differences in the position outputs of several friction force computation algorithms (e.g. position oscillations and steady drift), running on identical inputs (forces applied to the object), were examined in order to demonstrate better performance of suggested algorithm. Similarly, in [22], the authors provide “ground truth” data sets consisting of correlated input trajectories and output forces gathered in the real world; these datasets can then be used for an error-metrics-based evaluation of haptic rendering systems. Using root mean squared error metrics on the output force with regard to the real-world measurements, the study assessed the influence of friction compensation, force shading adaptation, and mesh resolution on the accuracy of haptic rendering. Although the former evaluation methods guarantee the resemblance of the interaction to the real-world behavior, they do not address the human perception of the simulated force feedback.

In many cases, high-quality haptic output is not required for realistic force perception. However, the perceived quality of haptic feedback received during object interactions in real versus virtual environments has seldom been addressed. In [21], the realism of virtual biological tissue-cutting using scissors was evaluated. Pre-recorded forces stemming from physical measurements and forces obtained from linear approximation of the empirical data were compared in a psychophysical study with experienced users. The participants were persons who were experienced with cutting the biological tissues being simulated by the display. Neither of the two methods tested was clearly rated as being less reliable. Given the high level of the participants’ expertise, the results indicate that both methods closely approximated real feedback.

Holistic evaluations of haptic systems which include comparisons to real-world references have only recently begun. One research group [17] performed a quantitative assessment of the fidelity of an open-loop haptic algorithm for rigid object interaction. Users were asked to tap on given samples while holding a PHANToM stylus and then to rate how realistically the simulation represented the experience of tapping on a real wood sample. In addition to the virtually-driven interaction, real samples of different materials were presented to the participants. Average realism ratings were evaluated using paired t-tests between all stimuli. Although the experience of tapping a real wooden sample could not fully be reproduced, the evaluation revealed a significant increase

in virtual realism when the novel open-loop haptic rendering technique was used.

In our previous work, we assessed a method for quantifying the fidelity of a haptic interaction with *deformable* objects. Two groups of soft objects were presented to users: six real silicone samples and six rendered objects. In a discrimination task, users were asked to determine whether the interaction was virtually driven via the haptic device or whether a real object was present. A signal detection measure, *d*-prime, was thus proposed as an evaluation metric [18].

However, discriminability may not be the best measure for evaluating haptic rendering fidelity. Another approach to perceptual validation, presented recently in [6, 5], is to use *similarity* as a criterion. Similarity has been proposed as a fundamental cognitive process underlying object recognition and categorization [11]. Using perceptual similarity between objects as a validation criterion has the advantage of encompassing higher-level, cognitive relationships between objects, which may be more closely correlated with realism and believability than a lower-level measure such as discriminability. Furthermore, MDS techniques can be used to analyze similarity data, yielding rich multidimensional stimulus representations. Thus, the purpose of the current study was to test whether such an approach could be usefully applied to the problem of haptic fidelity assessment.

3 VIRTUAL OBJECT REPRESENTATION

3.1 Force computation model

For simplicity and in order to maintain the broadest applicability, we selected a mass-spring model approach for our experiments [2]. The linear mass-spring model delivers fast, physically-based force computation for linear elastic bodies. While other deformation models could also have been used, many of them are computationally expensive and some limit the generality of interaction with 3D objects, which will be the target of our future work.

Our mass-spring model was based on the tetrahedral structure of the object we wished to simulate. The vertices of the tetrahedra were the mass points and the edges defined linear springs. Volume-preserving forces of the tetrahedral structure were imposed on each vertex. The dynamics of the mass-spring model was determined by Newton’s second law of motion:

$$M \frac{\partial^2 x}{\partial t^2} + D \frac{\partial x}{\partial t} + F_{internal} = F_{external},$$

where M is the mass matrix, D is the damping matrix, and F represents internal or external forces, respectively. To update the node positions and velocities, an explicit numerical integration scheme was used.

The coupling of the haptic device to the deformation simulation was accomplished using a point-based haptic proxy paradigm [23]. A linear spring between the proxy object and the real position of the tool tip was used to approximate the force applied by the user during the interaction. The external force generated by the haptic proxy model was distributed to the mass-spring model nodes of the surface triangle being contacted by the tool. The distribution was calculated using the barycentric coordinates of the interaction point with respect to the vertex positions of the contacted surface triangle. The deformation force computed from the proxy model was also filtered before it was rendered on the haptic device.

3.2 Haptic rendering conditions to be evaluated

In this study, we used MDS to compare two virtual haptic environments, which we shall refer to as *higher-fidelity* and *lower-fidelity* conditions. The two environments differed in terms of force and proxy-point position filtering. In the higher-fidelity environment,

we linearly blended the position of the proxy point within five steps. This interpolation was required due to a difference in refresh rates: in our system, the haptic loop runs at 1 kHz , while the simulation loop runs at 100 Hz . Although the proxy point was always locally adjusted within the haptic loop, large differences in its position appeared at each simulation step when the surface triangles of the deformed body moved. Interpolating between proxy point positions helped to mitigate this effect. Moreover, for enhanced stability we also filtered the last five values of the computed force vectors using a Gaussian filter before displaying the vectors on the haptic device.

In the lower-fidelity environment, we turned off the proxy point position blending. Furthermore, we simply averaged the last five samples of the force vector before rendering the force on the device. While the underlying deformation algorithm stayed the same, these changes led to a reduced fidelity of the haptic rendering. This manifested itself in minute force discontinuities and slightly diminished overall stability. Although this effect is difficult to describe, one can think of it as a reduction of the overall smoothness of the interaction. The resulting difference between the two conditions was not readily perceptible and could only be detected by experienced users of haptic interfaces.

3.3 Reference silicone samples

One of the main drawbacks of mass-spring models is having to set deformation parameters. Spring constants, masses, and mesh topology have to be adapted to obtain a specific deformation behavior. In order to achieve realistic behavior with our model, we determined the deformation parameters from reference silicone samples. To obtain a material with nearly linear-elastic properties, we used a two-part silicone rubber called ECOFLEX. We applied different mixtures of the Ecoflex 0030, Ecoflex 0040 and silicone thinner to prepare seven cylindrical ($80 \times 80\text{ mm}$) rubber phantoms for our studies. An aspiration test [20] was performed in order to determine the material properties of the silicone objects. According to these tests, the silicone phantoms could be considered as a neo-Hookean material with Young's moduli of 16, 20, 29, 38, 46, 69, 78 kPa , with Poisson's ratio assumed to be 0.499. There was a slight difference in height across the samples; the maximum difference was 2.5 mm .

3.4 Virtual samples

In addition to the real samples, we created seven virtual objects for our experiments. Each of the seven virtual samples was tuned according to the material properties of one of the seven manufactured silicone cylinders. The same deformation parameters were used for virtual models in both the lower and higher-fidelity rendering conditions. This study did not require that the stiffness of real and virtual objects be exactly matched. However, the virtual samples were designed to roughly cover the same stiffness range as the real samples as quantified by their Young's moduli. It should be noted that the overall force response of the virtual object is a result of several factors, including the mechanical model and the virtual coupling to the haptic device. For comparison, Figure 1 shows the force response of the virtual samples, which include all such factors, together with the force response of the real objects. Note that due to differences in the heights of the silicone samples, the force output measured on real samples starts at different height for each sample. The corresponding stiffness ratios of the stimuli, measured by a linear approximation of corresponding plots in Figure 1, and the height differences are listed in Table 1.

3.5 Parameter setting

For the mass-spring simulation, we created a mesh representation of a cylindrical object (80 mm high with a diameter of 80 mm). The object was uniformly tetrahedralized with 300 nodes, 1656 edges and 1156 tetrahedra. Afterwards, the deformation parameters of

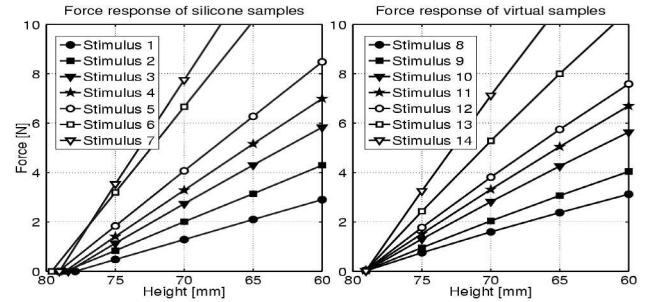


Figure 1: Force response of real samples (left), measured during an indentation experiment with 8 mm flat cylindrical indenter, and simulated force output of the virtual samples (right). The drift on the height axis for the real samples (left) indicates a slight variation in the samples' heights. The height of the virtual samples was calibrated to the height of the hardest real sample (stimulus 7).

Table 1: The stiffnesses of the real and virtual stimuli, measured by a linear approximation of corresponding plots in Figure 1, and height differences of the real samples in comparison to the fixed height of the virtual samples (positive values indicate shorter samples).

Real stimuli	1	2	3	4	5	6	7
Stiffness [10^{-2} N/mm]	16	23	32	37	44	69	84
Height difference [mm]	1.1	0.4	0.6	0.4	0.0	-0.5	0.0
Virtual stimuli	8	9	10	11	12	13	14
Stiffness [10^{-2} N/mm]	16	21	29	34	38	55	72

the mass-spring system for each virtual object were tuned to approximate one of the assembled silicone cylinders. The mass of the reference samples (400 g) was distributed to nodes of the virtual sample according to [9]. Mass moments up to the second order were matched to the moments of a cylindrical object according to

$$\int_{\text{Body}} u_x^j u_y^k u_z^l \rho(\vec{u}) d\vec{u} = \sum_{i=1}^n u_{xi}^j u_{yi}^k u_{zi}^l m_i, j+k+l \leq 2$$

where the mass of each node is m_i and its position is u_i . To solve this under-determined system, we applied a least-squares minimization of the distributed masses to the masses m_i of each node v_i estimated by the volume of tetrahedra incident upon v_i :

$$\min \left(\sum_i (m_i - c \sum_{v_i \in T} \text{Volume}(T))^2 \right),$$

where c is a mass correcting constant and T is a tetrahedron defined by the model mesh. The node masses in the virtual model varied from 0.4 g to 3.3 g . The stiffness of each spring e was set proportional to the Young's modulus E , based on [13]

$$k_e = \frac{E \sum_{e \in T} \text{Volume}(T)}{\text{Length}(e)^2}.$$

The simulation time step was set according to [8] using:

$$\Delta t = \sqrt{\frac{\text{mass}}{n \pi^2 k_{\max}}},$$

where k_{\max} is the maximum spring stiffness, n is the number of nodes, and mass is the total mass of the object. This equation estimates the time step beyond which the system of equations of motion is divergent when using an explicit numerical integration method. Node and spring damping parameters were manually selected to ensure that the dynamic deformation remained stable.

Finally, the constants for the virtual coupling to the proxy object were selected. This connection should be as stiff as possible

while still allowing stable interaction. In general, a stiffness proportional to that of the virtual samples was used (values varied from 0.7 N/mm to 1.5 N/mm).

4 SYSTEM SETUP

We chose SensAble’s PHANToM device to provide the interface for our virtual environment. This decision was made on the basis of availability. Due to the limitations in the device’s force response [4, 3], the PHANToM is not the optimal choice for faithfully displaying the complete stiffness range of our virtual objects. Nevertheless, with some adaptations of hardware and software (described below), we were able to provide adequate feedback in all cases. With regard to this, we wish to emphasize that our analysis did not focus on evaluating any specific device but rather on evaluating the fidelity of the complete haptic interaction environment.

The choice of the PHANToM haptic device introduced limitations in maximal force output. From empirical tests, this was estimated to be around 4.2 N for a long-term peak force. If this value is exceeded, force feedback is switched off and the application driving the PHANToM stops. To keep the application running, we limited the computed force to the estimated long-term peak force. The maximum value of 4.2 N guarantees realistic force response for indentations up to 6 mm for the hardest sample. For penetrations deeper than 6 mm , the force output was held constant. To increase the interaction depth (i.e., by increasing the maximum force output of the haptic device), an additional weight of 200 g was attached to the other end of the PHANToM arm, directly on one of its motors. This weight constantly pulled the PHANToM stylus upward, which provided an additional increase of the maximum force by about 2 N in the upward direction. For our experiment, we could therefore render forces up to 6.5 N , which was sufficient for indentations in the hardest sample of up to 9 mm . Table 2 shows the maximum allowable indentation depths for each virtual object.

Table 2: The Young’s modulus parameter of the virtual samples (stimuli 8 – 14) and maximum depth of realistic interaction. At greater depths, rendered forces would have exceeded the maximum force limit of the haptic device and were therefore held constant at 6.5 N .

Virtual stimuli	8	9	10	11	12	13	14
Young’s modulus [kPa]	16	20	29	38	46	69	78
Maximum depth [mm]	43	32	22	19	17	11	9

Attaching an additional weight to the PHANToM arm required real-time, software-driven gravity compensation. For both interactions with real and virtual objects, a force which counterbalanced the additional weight was generated on the haptic device throughout the whole experiment. When a virtual object was being presented, this force was added to the forces stemming from the computation model of the virtual sample. The gravity compensation force was played back from a pre-recorded force profile. Forces sufficient to maintain the position of the PHANToM stylus were recorded at 23 positions spread vertically across the interaction axis and stored as the gravity compensation force profile.

The test application was run on a Linux PC with 2 P4 2.8GHz processors. We used two asynchronous threads: one for haptic force rendering using the proxy model and one for deformation calculation. During the experiments, the haptic thread ran at 1 kHz , while the deformation calculation ran at approximately 100 Hz .

5 SIMILARITY-BASED VALIDATION EXPERIMENT

We evaluated the two haptic rendering conditions described in Section 3.2 within an indentation experiment, in which the users prodded samples with a metal rod and rated the similarity between pairs

of objects drawn from a set of seven real and seven virtual objects. The proximity data were analyzed using multidimensional scaling as described in the following section. We hypothesized that a single perceptual dimension corresponding to stiffness would be sufficient to explain similarity data in the higher-fidelity rendering condition, whereas an additional dimension corresponding to the distinction between real and virtual stimuli would be required in the lower-fidelity condition.

5.1 Participants

Ten right-handed participants (19 – 43 years old) were paid 32 EUR to perform the experiment. None of them reported any haptic deficit due to accident or illness. None of the participants had used a kinesthetic haptic device before and all of the users were naive about our research in the haptic field.

5.2 Experimental stimuli: Real and virtual objects

Participants performed similarity ratings amongst a set of fourteen objects. This set consisted of seven silicone cylinders and seven virtual cylinders which varied in stiffness. The same set of real and virtual objects was used for the whole experiment. However, half of the subjects experienced the virtual objects under the higher-fidelity filtering condition, while the other half experienced the virtual objects under the lower-fidelity filtering conditions.

Importantly, participants were not told about the differences amongst the objects being presented (real/virtual, soft/hard). The real samples were not visible to the participants during the experiment; they were hidden under a cardboard box as shown in Figure 2. On each trial, the rod slid through a small hole in the box, presenting the same visual feedback to the participant. The PHANToM device, attached to the indenter, was not hidden from the users. Participants, who asked about it were told that it was used to record indenter positions, velocities, and applied forces during the interaction.

The experiment was restricted to the study of one-dimensional interactions with soft objects. This choice was made to minimize the influence of contact effects such as friction, multiple contact points, torques, etc. Prodding of the soft bodies was only possible on the top surface of the cylinder, around the middle of the top surface, and only in the vertical direction.

5.3 Experimental apparatus

The experiment was performed with a PHANToM haptic device, with a 250 mm long pen-shaped rod. One end of the pen was attached to the PHANToM robot arm (Figure 2). A flat, cylindrical indenter (radius of 4 mm) was fastened to the other end of the pen. The flat indenter was chosen to alleviate the need to compensate for the indenter’s shape in the force response algorithm. The cylinders’ size was set to ensure sufficient penetration depth and elastic deformation of the real samples. We instructed participants to penetrate to a maximum depth of $20 - 25\text{ mm}$. A mark was made on the stylus pen and participants were told that it should be kept above the construction’s keyhole. However, no rigid stop prevented the users from pushing deeper into the object. Participants were asked to get used to the interaction depth during the training session, however during the experiment, they were encouraged not to watch the mark.

The stylus was held upright by three cylinders which ensured that the indenter moved smoothly along a vertical axis. To avoid sideways vibrations due to stiff contact between the stylus and the cylinders, the computed PHANToM forces were projected along the vertical direction. Beneath the indenter, seven real samples were placed on a rotatable support, while leaving room for virtual samples. All the samples, and the tip of the indenter, were covered by

a cardboard box upon which the participant could lean his/her arm during the experiment (see Figure 2).




Figure 2: Experimental setup exposed (left) and covered for the experiment (right). The setup consisted of a PHANTOM, an indenter attached to the PHANTOM, three cylinders to prevent sideways motion, and a rotatable plate for the silicone samples.

To mask the sound of the PHANTOM motors, users were asked to wear headphones as acoustic ear protectors. No visual feedback, e.g. from the virtual simulation, was provided to the users. The operator was on the other side of the table, monitoring the experiment and changing the samples beneath the indenter. The experimenter's motions were kept as consistent as possible across trials and were the same regardless of whether the target object was real or virtual. Since the deformation model used was optimized for quasi-static interaction, participants were instructed to limit their indentations to slow movements. This resulted in contact velocities of about 100 mm/s on average. The rationale for this limitation was to avoid strong impacts, which would in turn cause some complex dynamic effects that might not be accurately reproducible with the mechanical model.

5.4 Experimental procedure

The experiment consisted of two phases: a training session and a testing session. In both phases, the participants' task was the same: they prodded two objects, one after another, and then rated the similarity between the two objects on a scale between one (low similarity) and seven (high similarity) by typing the number on the keyboard. In the training phase, participants practiced using the setup and interacting with the objects until they were able to prod the objects as directed. Participants held the metal rod in their right hand between their thumb and index finger. They were instructed to slowly approach the objects with the indenter, to hold the indenter in their hand during the interaction, and to respect the maximum indentation depth of 20 mm . All of the objects were introduced to the participants so that they could indent them, notice differences, and establish their rating scale. Participants were told that similarity judgments should not take the slight difference in sample height (max. 2.5 mm) into account. In the training phase, every participant went through the same sequence of presented stimuli. The training phase lasted for about 20 minutes, during which it was possible to perform at least 50 trials. At the end of the training session, all participants felt comfortable with the task and the setup, and were ready to run the main test.

The testing phase consisted of five blocks of similarity ratings. In each of the five blocks, all pairwise combinations of the fourteen objects (105 combinations) were presented to the participant in random order. On each trial, the experimenter set the first object below the indenter by turning a plate with all samples to the right position (Section 5.3) or by altering appropriate settings for the virtual objects, and then asked the user to test the sample. The user was allowed to prod the sample just once. After the user lifted the indenter, the experimenter presented the second object, and again asked

the user to prod the sample and provide the similarity rating. The testing phase lasted about 3 h. After completing the experiment, participants were asked to fill in a questionnaire in which they explained how they performed the similarity ratings, and described the samples used in the test.

5.5 Multidimensional scaling analysis

Subjects' similarity ratings (values ranging from 1 to 7) were converted to dissimilarities (values ranging from 0 to 1) and the average pairwise dissimilarities were computed over all subjects for each rendering condition. The mean dissimilarity data were analyzed separately for each condition using the non-metric MDS algorithm (MDSSCALE) implemented in MATLAB. As opposed to the classic metric MDS techniques, which fit exact proximity values, non-metric MDS techniques only take the *rank-order* of the pair-wise proximity values into account and are thus better suited for human similarity data. The algorithm searches for an output configuration in a space of fixed dimensionality and returns a goodness-of-fit measure, Kruskal's S1 Stress, which is the normalized difference between the fitted distances in the output configuration and the observed proximities [7]. Stress values are used to determine the required dimensionality of the output configuration. Given our set size of 14 objects, stress values below 0.2 can be taken as an indication that the dimensionality of the output space is sufficient to faithfully represent the input proximities [1].

Applying MDS analysis on similarity data in this setting provides a quantitative method for identifying perceptual differences between real and virtual haptic objects, as well as between objects rendered using the two different rendering methods. Critically, the object sets we used included both real and virtual samples which were expected to vary in stiffness. Thus, for each of the rendering methods, we used MDS to test whether subjects' similarity ratings could be explained by a unidimensional configuration, where the single dimension corresponded to stiffness, or whether additional perceptual dimensions (e.g., a systematic difference between real and virtual stimuli) are necessary to explain the similarity data. In addition to the question of overall dimensionality, we were interested in visualizing output configurations in order to better understand the perceptual differences between individual samples.

6 RESULTS

Similarity data from each of the participants was grouped according to the rendering condition under which they experienced the haptic objects and the average similarity was taken over all participants in each group. Although patterns of changing similarity could be observed in this data, performing MDS analysis of the similarity data provides a richer perspective including the following kinds of information: how many dimensions of variation in the stimuli were apparent to the participants; whether one of these dimensions corresponded to the material property being simulated (stiffness, in this case); whether one or more undesirable perceptual dimensions were also apparent to the subjects (e.g., a dimension along which real and virtual stimuli are separated); the relative weights of the desired dimensions relative to undesired perceptual dimensions; interstimulus distances in the perceptual space.

6.1 Perceptual dimensions of MDS configurations

The first question addressed by MDS analysis is that of the appropriate dimensionality of the output configuration. This is classically determined by inspection of the MDS stress plot, as described in Section 5.5. Figure 3 shows the stress plots for MDS solutions with dimensionalities of 1 – 5 obtained using mean similarity from the higher-fidelity and the lower-fidelity condition. In the former,

stress drops below the 0.2 threshold for a one-dimensional solution, indicating that one perceptual dimension is sufficient to explain the similarity data. In contrast to this, for the lower-fidelity condition, stress was 0.25 for a one-dimensional solution, and drops sharply to 0.05 for a two-dimensional one. This clearly indicates that two perceptual dimensions are required to explain this similarity data.




Figure 3: Stress as a function of the dimensionality of the MDS solution computed using similarity data from each of the two rendering conditions. A dashed line is drawn at a stress level of 0.2; configurations with stress below 0.2 are considered to adequately explain similarity data.

MDS does not provide labels for the dimensions of the output configuration; these must be interpreted by visual inspection. The output configurations generated by two-dimensional MDS solutions are shown in Figure 4 for the higher-fidelity (left) and the lower-fidelity condition (right). Note that for ease of comparison, the configurations shown here have been scaled and rotated using the Procrustes transform implemented in MATLAB such that the points representing the real objects were brought as close as possible to equidistant points along the x-axis. For both conditions, the first perceptual dimension recovered by MDS (plotted along the x-axis) clearly corresponded to stiffness.

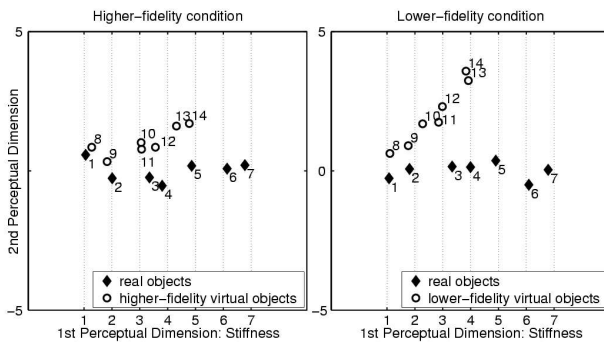


Figure 4: Two-dimensional MDS maps for the higher fidelity condition (left) and lower fidelity condition (right).

As mentioned above, in the higher-fidelity condition, the single dimension of stiffness sufficed to explain the similarity data. This finding provides support for the perceptual validity of the higher-fidelity environment, insofar as differences between real and virtual objects did not play a significant role in determining perceptual similarity. In contrast to this, a second perceptual dimension *is* required to explain similarity data in the lower-fidelity condition. From the lower-fidelity configuration, we see that this dimension causes a perceptual separation between the real and virtual samples, and that the separation increases with stiffness. Note that a similar, albeit much smaller trend is also visible in the higher-fidelity condition. This trend increases at higher stiffness levels. Comparing the results from the two conditions, it appears that the lower-fidelity condition exaggerated or compounded a difference between real and virtual stimuli which was also present in the higher-fidelity condition, but only had a negligible perceptual effect.

6.2 Topology of MDS configurations

The MDS maps also provide more detailed topological information about stimulus ordering and clustering in perceptual space. First, participants were able to recover the rank order of real samples along the stiffness dimension in both rendering conditions. In addition, the spacing between real samples is similar under both conditions, indicating that subjects' perception of the real samples' stiffness was not influenced by the rendering mode of virtual samples. Participants also recovered the rank order of virtual stimuli to a large extent, although stimuli 10 and 11 were perceived as almost equally stiff in the high-fidelity condition, while stimulus 14 was perceived as slightly less stiff than stimulus 13 in the lower-fidelity condition. However, differences in stiffness levels were not equal between all samples. For the real samples, there was a greater perceptual difference in stiffness between samples 2 and 3 and between samples 5 and 6 than between other neighbouring pairs. These differences were constant across rendering conditions. Virtual stimuli cluster in roughly the same way, yielding a group of lower stiffness objects (real: 1, 2, virtual: 8, 9), a group of moderate stiffness objects (real: 3 – 5, virtual: 10 – 12), and a group of higher stiffness objects (real: 6, 7, virtual: 13, 14).

Interestingly, this clustering pattern is according to the one predicted by the Young's modulus of the samples, i.e., the spacing between the samples in the map matches the differences between the object stiffnesses quite well. However, the absolute range of perceived stiffness is compressed for the virtual objects compared to the real ones. In addition, the *degree* of compression is different in the two rendering conditions: the stiffest virtual objects (13, 14) have a perceived stiffness equivalent to that of real object 4 in the lower-fidelity condition, while their stiffness is closer to that of real object 5 in the higher-fidelity condition.

6.3 Results of debriefing questionnaires

In the questionnaires, participants were asked to describe the objects and to explain how they performed similarity ratings. Interestingly, none of the participants detected the presence of real and virtual objects in the study. All but one assumed that all samples were real objects; the other participant thought that all samples were virtual. His assumption was likely due to the fact that he was aware of the existence of haptic device technology and assumed that all stimuli would be presented using the device.

In their descriptions of the objects, participants mentioned several properties which they had used for similarity judgements, most of which could be related to stiffness. For example, participants described objects as being soft and homogeneous and likened them to foam, rubber, jelly, sponge, or metal springs. The fact that stiffness-related terms were mentioned so often confirms that the first dimension in the MDS maps can be interpreted as stiffness.

A smaller number of participants described objects as having a rigid skin or shell, with a softer interior. This was sometimes referred to as *hollowness*. One interpretation of this could be that it was just a different description of the stiffness dimension. However, it could also be possible that this relates to a second, unexpected perceptual dimension in the experiment. The hollowness attribute might have been caused by the maximum force limitation of the haptic device. When large penetration depths were used during interaction, the displayed force had to be held constant, which led to a sudden pop into the sample and which may have induced a sense of "hollowness". In order to rule out the possibility that this could correspond to the second dimension we found in the MDS plots, we grouped the participants according to the penetration depths they used during interactions and performed another MDS analysis of this data. However, the results showed that penetration depth did not cause the separation of objects along the second dimension.

Nevertheless, participants did not mention differences in haptic rendering fidelity in the questionnaires. These effects were of course quite subtle and could not easily be detected, but it could have been expected that for the lower-fidelity condition users would report another criteria which they used for their similarity assessment. Clearly, further work is needed in this direction, especially requiring a larger number of participants.

6.4 Summary

The MDS analysis revealed that participants were able to recover stiffness changes across both real and virtual objects, in both lower-fidelity and higher-fidelity environments. We found a noticeable perceptual difference between real and virtual objects in the lower-fidelity condition, especially at high stiffness levels. In contrast to this, in the higher-fidelity environment, real and virtual objects were quite similar at low stiffness levels.

Although real and virtual objects showed a trend towards differentiation along a second dimension with increasing stiffness, this dimension was found to be perceptually more significant in the lower-fidelity condition. Plotting the objects in a two-dimensional perceptual space also revealed that stiffness differences roughly correlated with differences in Young's moduli. Finally, the verbal reports allowed us to confirm that stiffness was, indeed, the most important perceptual dimension in both rendering conditions.

7 CONCLUSION AND OUTLOOK

In this paper, we have shown how MDS analysis of human similarity judgments on a set of real and virtual objects can be used to gain both quantitative and qualitative insight into the relationships between haptic rendering parameters and human perception. In this specific test, two haptic rendering conditions were compared. MDS analysis revealed a clear perceptual distinction between real and virtual objects in the lower-fidelity condition, but not in the higher-fidelity condition. In addition, MDS maps of the objects in human perceptual space showed that real and virtual objects are most similar at lower levels of stiffness and become more distinguishable as stiffness increases. Removing smoothing of proxy-point positions and rendered forces compounded this effect.

This preliminary study demonstrates a new application of MDS techniques, showing how they can be used to validate a particular rendering algorithm using human perception of real objects as a benchmark. The method can also be used to compare different rendering approaches; in the present study, this was done indirectly by having two groups of participants, each of which compared real samples against virtual samples rendered with different feedback fidelity. Another possibility would be to pit two or more rendering algorithms directly against one another in a single condition.

Future work will extend the work presented here to investigate a wider range of rendering parameters and interaction styles (i.e. allowing for fast and oscillatory movements, and strong impacts in order to investigate the dynamic properties of the simulated bodies). Alternative haptic devices will be considered with the aim of avoiding the hardware limitations encountered with the PHANTOM. The role of interaction velocity and/or penetration depth on similarity ratings need to be systematically tested using a larger number of participants. Finally, in future experiments, it will also be of interest to further quantify verbal responses (e.g., by scoring questionnaires) in order to provide more quantitative correlations between MDS dimensions and conscious human perception. Together, these studies should provide valuable insight into haptic perception of virtual deformable objects.

ACKNOWLEDGEMENTS

The authors would like to thank to Gustavo Lunardi, Marc Hollenstein and

Alessandro Nava for measuring the parameters of the real silicone samples, and Christian Wallraven for helpful insights and support. This research was supported by the TOUCH-HapSys project EU CEC IST-2001-38040.

REFERENCES

- [1] I. Borg and P. Groenen. *Modern Multidimensional Scaling: Theory and Applications*. Springer, 2nd edition, 2005.
- [2] D. Bourguignon and M.P. Cani. Controlling anisotropy in mass-spring systems. In *11th Eurographics Workshop Proc.*, pages 113–123, 2000.
- [3] G. Campion and V. Hayward. Fundamental limits in the rendering of virtual haptic textures. In *WHC'05*, pages 263–270, 2005.
- [4] M.C. Cavusoglu, D. Feygin, and F. Tendick. A critical study of the mechanical and electrical properties of the PHANTOM haptic interface and improvements for high performance control. *Presence*, 11(6):555–568, 2002.
- [5] T. Cooke, F. Jäkel, et al. Multimodal similarity and categorization of novel, three-dimensional objects. *Neuropsychologia*, march 2006.
- [6] T. Cooke, F. Steinke, et al. A similarity-based approach to perceptual feature validation. pages 59 – 66, NY, USA, 08 2005. ACM Press.
- [7] T.F. Cox and M.A. Cox. *Multidimensional Scaling*. Chapman & Hall, 2nd edition, 2001.
- [8] H. Delingette. Towards realistic soft tissue modeling in medical simulation. In *Proceedings of the IEEE : Special Issue on Surgery Simulation*, pages 512–523, April 1998.
- [9] O. Deussen, L. Kobbelt, and P. Tücke. Using simulated annealing to obtain good nodal approximations of deformable bodies. In *Sixth Eurographics Workshop on Simulation and Animation*, 1995.
- [10] P. Dupont, B. Armstrong, and V. Hayward. Elasto-plastic friction model: Contact compliance and stiction. In *Proc. 2000 American Control Conference*, pages 1072–1077, Chicago, Ill., USA, June 2000.
- [11] S. Edelman. *Representation and recognition in vision*. MIT Press, 1999.
- [12] B. Forsyth and K.E. MacLean. Predictive haptic guidance: Intelligent user assistance for the control of dynamic tasks. In *IEEE Trans.*, volume 12, pages 103–113, January 2006.
- [13] A.V. Gelder. Approximate simulation of elastic membranes by triangulated spring meshes. *Journal of Graphics Tools*, 3(2):21–42, 1998.
- [14] V. Hayward and O.R. Astley. Performance measures for haptic interfaces. In *In Robotics Research: The 7th International Symposium*, pages 195–207. Springer Verlag, 1996.
- [15] M. Hollins, R. Faldowski, et al. Perceptual dimensions of tactile surface texture: A multidimensional scaling analysis. *Perceptual Psychophysics*, 54(6):687–705, 1993.
- [16] S.H. Hsu, W. Chang, and M. Chuang. Effects of geometric form features on three-dimensional object categorization. *Perceptual Motor Skills*, 100(3 Pt 2):899–912, 2005.
- [17] K.J. Kuchenbecker, J. Fiene, and G. Niemeyer. Event-based haptics and acceleration matching: Portraying and assessing the realism of contact. In *World Haptics Conference*, Pisa, IT, March 2005.
- [18] P. Leškovský, M. Harders, and G. Székely. Assessing the fidelity of the haptically rendered deformable objects. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, IEEE VR*, 2006.
- [19] R.D. Melara and D.J. Day. Primacy of dimensions in vibrotactile perception: an evaluation of early holistic models. *Perceptual Psychophysics*, 52(1):1–17, 1992.
- [20] A. Nava, E. Mazza, et al. Determination of the mechanical properties of soft human tissues through aspiration experiments. In *MICCAI 2003*, volume 2878 of *LNCS*, pages 222–229, 2003.
- [21] A.M. Okamura, R.J. Webster, et al. The haptic scissors: cutting in virtual environments. In *IEEE ICRA*, pages 828–833, 2003.
- [22] E. Ruffaldi, D. Morris, et al. Standardized evaluation of haptic rendering systems. In *Haptic Interfaces for Virtual Environment and Teleoperator Systems, IEEE VR*, 2006.
- [23] D.C. Ruspini, K. Kolarov, and O. Khatib. The haptic display of complex graphical environments. In *Computer Graphics (SIGGRAPH 97 Conference Proceedings)*, pages 345–352, 1997.
- [24] R.N. Shepard. Analysis of proximities as a technique for the study of information processing in man. *Human Factors*, 5:33–48, 1963.