Acta Psychologica 132 (2009) 259–266

Contents lists available at ScienceDirect

Acta Psychologica

journal homepage: www.elsevier.com/locate/actpsy

Perceptual equivalence between vision and touch is complexity dependent ☆

F. Phillips [a,*], E.J.L. Egan [b], B.N. Perry [c]

[a] Skidmore College, Psychology and Neuroscience, Saratoga Springs, NY 12866, USA
[b] The Ohio State University, Psychology, Columbus, OH, USA
[c] Burnt Hills-Ballston Lake HS, Burnt Hills, NY, USA

ARTICLE INFO

Article history:
Received 25 February 2009
Received in revised form 9 July 2009
Accepted 16 July 2009
Available online 18 August 2009

PsycINFO classification:
2320

Keywords:
Touch
Vision
Shape
Psychophysics

ABSTRACT

We experience the shape of objects in our world largely by way of our vision and touch but the availability and integration of information between the senses remains an open question. The research presented in this article examines the effect of stimulus complexity on visual, haptic and crossmodal discrimination. Using sculpted three-dimensional objects whose features vary systematically, we perform a series of three experiments to determine perceptual equivalence as a function of complexity. Two unimodal experiments – vision and touch-only, and one crossmodal experiment investigating the availability of information across the senses, were performed. We find that, for the class of stimuli used, subjects were able to visually discriminate them reliably across the entire range of complexity, while the experiments involving haptic information show a marked decrease in performance as the objects become more complex. Performance in the crossmodal condition appears to be constrained by the limits of the subjects' haptic representation, but the combination of the two sources of information is of some benefit over vision alone when comparing the simpler, low-frequency stimuli. This result shows that there is crossmodal transfer, and therefore perceptual equivalency, but that this transfer is limited by the object's complexity.

1. Introduction

The problem of visual-haptic sensory integration has a long history, perhaps best posed in the notorious philosophical puzzle known as the *Molyneux problem* (Locke, 1694; Locke, 1978). In 1690, William Molyneux asked of John Locke – if a blind man were given sight, would he be able to recognize objects he had previously only touched? While there is ground-breaking research currently being conducted on this specific query (Ostrovsky, Andalman, & Sinha, 2006), the focus of this paper is based on a basic assumption Molyneux made – That it is possible for humans born *with* vision to transfer haptic information and prior knowledge into a representation that is visually salient and vice versa. To some degree this task is trivial for normally sighted individuals, which is probably why it was overlooked for so long. Gibson (1966) and his student Caviness (1962), Caviness (1964) were some of the first to tackle this question empirically. They found that, with practice, subjects could easily complete a series of crossmodal matching tasks that used hand-sized sculpted stimuli. Unfortunately, the bulk of these results went unpublished and their published allusions to this work merely mention that some perceptual equivalence exists.

The aim of the current research is to examine how three-dimensional (3D) object features affect this perceptual equivalence. A 3D shape can be decomposed at various levels using any of a variety of methods. For example, if a shape is complex, it can be broken down into parts using various rules (Biederman, 1987; De Winter & Wagemans, 2006; Hoffman & Richards, 1984; Hoffman & Singh, 1997; Norman, Phillips, & Ross, 2001). Alternatively, if a shape is a simple surface a geometrical measurement is all that is necessary to fully represent it. For example, Kappers, Koenderink, and Oudenaarden (1997) and Vogels, Kappers, and Koenderink (1997) tested the perceptual equivalence of the two perceptual systems for surfaces with a single shape characteristic. When subjects match parabolically curved surfaces with visually presented cross-sections, judgements are ordinally correct but are systematically scaled. It is difficult though to generalize from such simple geometric surfaces to more globally complex 3D shapes. In an extensive series of free-sorting experiments Garbin and Bernstein (1984), Garbin (1988), Garbin (1990) employed hand-sculpted stimuli similar to those used by Caviness (1962), Caviness (1964) and Gibson (1966). They developed a quantitative way to measure perceptual

☆ Experiments were designed by FP and EJLE, implemented by BP and FP, analyzed by FP and the paper was written by EJLE and FP.

* Corresponding author. Tel.: +1 518 580 5313; fax: +1 518 580 5319.
E-mail addresses: flip@skidmore.edu, flip@me.com (F. Phillips), egan.51@osu.edu (E.J.L. Egan).
URL: http://www.skidmore.edu/flip (F. Phillips).

0001-6918/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.actpsy.2009.07.010

equivalency based on multidimensional scaling. The results show that stimuli that were sorted similarly by each modality independently were easier to identify in a crossmodal matching task. Their results, however, were unable to disambiguate which shape attributes were responsible for the decreases in perceptual equivalency.

As the spatial frequency of 3D shape characteristics increases, unique low frequency structures transition from a macroscale, through an intermediate mesoscale, and eventually become microscale texture structures (Phillips, Todd, Koenderink, & Kappers, 2003; Phillips, Thompson, & Voshell, 2004). The macro- and mesoscales are roughly equivalent to the more broadly considered 'global' and 'local' characteristics of a shape. It has been suggested that, at very high feature frequency, perceptual equivalence is high (Picard, 2006; Tiest & Kappers, 2006; Tiest & Kappers, 2007). Picard (2006) was able to replicate Garbin's results using similar techniques but with textures as compared to global 3D shapes. At this end of the spatial frequency dimension haptic perception is superior to vision. This is counterevidence to a longstanding belief that vision is more reliable and therefore superior (Gibson, 1962; Rock & Harris, 1967; Rock & Victor, 1964). If prisms are used to create the visual illusion of seeing a flat surface as curved, our haptic perceptions may also wrongly conclude that the surface is flat (Gibson, 1966). The current trend of empirical evidence suggests that sensory superiority is stimulus dependant (Heller & Clyburn, 1993; Lakatos & Marks, 1999; Navon, 1977). The haptic system has a propensity to perform better on local compared to global-shaped perception, while the opposite is true for vision. Furthermore, it has been proposed that sensory integration occurs in a statistically optimal manner (Ernst & Bülthoff, 2004). It is clear that the information content of the stimuli used in a particular experiment can matter greatly.

Norman, Norman, Clayton, Lianekhammy, and Zielke (2004) confirmed this result with a recent series of experiments using a stimulus set that encompassed a single natural object category – bell peppers. The unimodal visual and then haptic discriminations were the easiest while crossmodal discriminations were least successful. Based on these results, they made the keen observation that performance differed greatly depending on specific stimuli pairings despite the stimuli sharing similar basic characteristics. They have since extended these experiments to learning (Norman, Clayton, Norman, & Crabtree, 2008) and aging (Norman et al., 2006; Norman, Bartholomew, & Burton, 2008). However, in all of their studies, only qualitative inferences could be made as to why there were discrimination problems with specific stimuli since no meaningful quantitative measurements were derived from the shapes of the bell peppers.

Within a range of global shapes, the dimension of shape complexity has been shown to effect visual shape recognition for 2D (Kayaert & Wagemans, 2009) as well as 3D shapes (Op de Beeck, Torfs, & Wagemans, 2008). Complex visual stimuli, which can be defined as having a large amount of high spatial frequency geometric components, are harder to identify in matching tasks compared to simpler or symmetric stimuli. Shape similarity even interacts with where information is computed within the brain according to fMRI data (Op de Beeck et al., 2008). Many brain regions that were originally thought to be modality-specific are activated by both visual and haptic stimulations. This is supported by studies of brain-damaged patients (Feinberg, 1986; James, 2003; Lacey, Campbell, & Sathian, 2007), single-cell recording studies (Bruce, Desimone, & Gross, 1981; Duhamel, Colby, & Goldberg, 1998), other imaging studies (Amedi, Malach, Hendler, Peled, & Zohary, 2001; Blake, Sobel, & James, 2004; Macaluso & Driver, 2001; Sathian, Zangaladze, Hoffman, & Grafton, 1997), and transcranial magnetic stimulation (Zangaladze, Epstein, Grafton, & Sathian, 1999).

The aim of the current research is to investigate how 3D shape complexity will affect visual and haptic perceptual ability as well as perceptual equivalence. Our stimuli are similar to those used by Gibson, Garbin and Norman, in that they are globally complex hand-sized shapes with sculptural qualities. Our stimuli have the advantage of being defined along a dimension of complexity defined by the spatial frequency of 'bumpy' geometric features. To investigate the relationships between haptic and visual information, it is crucial to establish baseline performance of each modality by considering the senses separately. Our first two experiments consider haptics and vision, respectively, and our final experiment examines their interaction.

2. Experiment 1

We begin by performing a haptic discrimination experiment using a set of systematically varying stimuli.

2.1. Method

In this experiment, subjects performed a simple discrimination using a single interval two-alternative forced choice (2AFC) paradigm. Two sculpted stimuli were presented haptically, without visual information, and subjects were instructed to indicate if the presented shapes were geometrically the same or different.

2.1.1. Stimuli

The full set of stimuli used in this and subsequent experiments are shown in Fig. 1. The set consist of 25 globally convex, natural appearing, noisy objects created using the techniques outlined in our previous work (Phillips, 2004). Basically, these objects are spheres subjected to a series of parameterized, pseudo-random transformations at various scales. In the computer graphics literature this technique is sometimes known colloquially as 'turbulent noise'. The objects' global shapes differ systematically along two dimensions: those of spatial frequency and amplitude. Variations of these two shape parameters subsequently give rise to objects whose shapes differ in overall visual complexity yet remain phenomenally similar throughout the set.

The resulting stimuli are ordered by spatial frequency – roughly the number of bumps and dimples per object, labeled from 1 to 25 starting with the smallest spatial frequency. The amplitudes, e.g., the height and depth of the bumps and dimples, decrease as a function of frequency in order to normalize the shapes' volume.

The underlying noise used to generate the objects is of a constant 'shape' for all stimuli – as frequency is increased the sampling range of the noise signal increases. This means that the objects share an underlying self-similarity due to this noise constancy but that similarity occurs at different scales across the object range. For example, note the similarity of objects 11 & 12, 14 & 15, and others in Fig. 1.

The resulting objects have an average diameter of ≈ 50 mm, all fit within a $55 \times 55 \times 55$ mm bounding box, and range from relatively smooth and free of notable inflections to a maximum of ca. 10–12 bumps per object. This size was chosen to allow the objects to be easily manipulated by hand, without requiring the traversal of large surface areas. Thus, through active touch, aspects of the global shape could be derived all at once, if desired, rather than in a piecewise fashion.

Each object was printed in ABS plastic using a Dimensions rapid prototyping three-dimensional printer (Stratasys, Inc.). Molds were taken from each and plaster casts were made to facilitate convenient reproduction of the stimulus set.

The resulting cast objects were painted a saturated blue. Of course, in this experiment the color and shading characteristics are of no importance since the task is completed totally by touch,

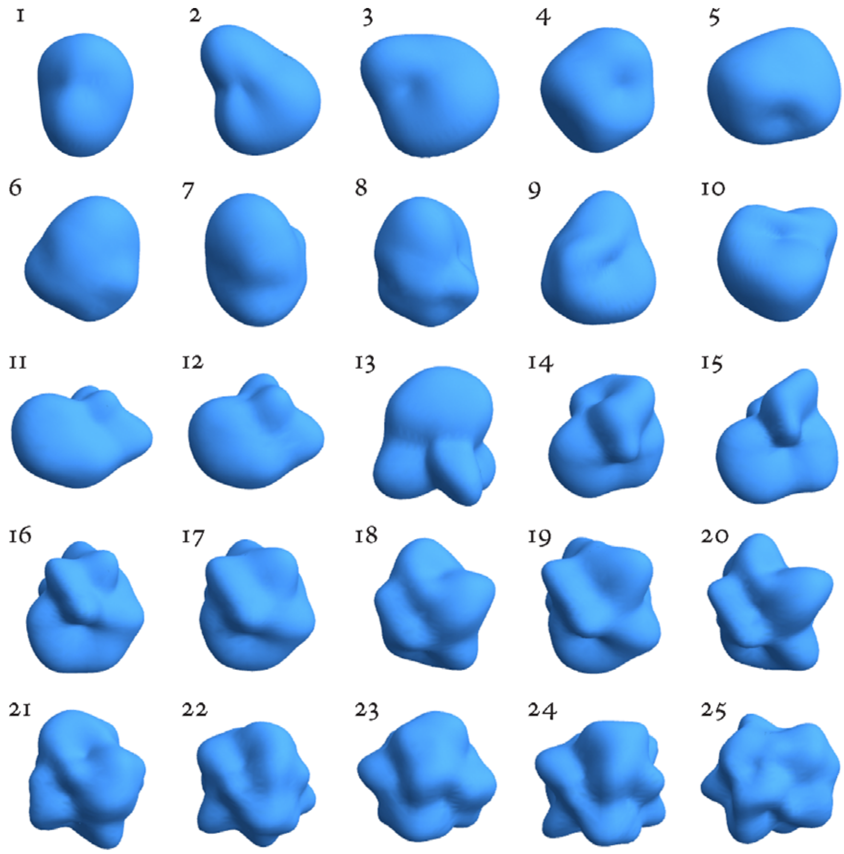

Fig. 1. The entire stimulus set consisting of 25 globally convex, natural appearing, noisy objects, ordered by spatial frequency.

but the paint provides a smooth even surface on the stimuli to prevent intentional or inadvertent use of potentially diagnostic local features in the discrimination task.

2.1.2. Procedure

Subjects were seated at a table and instructed to adjust their chair such that they were situated at a comfortable height. A wooden shelf and black cloth served to block the view of the stimuli from the subject. The subject placed their hands under the shelf, through the black cloth for the duration of the experiment. The experimenter sat across from the subject, fully obscured from view, with two sets of objects laid out in trays. Each object was marked with a UV-sensitive ink, indicating the stimulus number. A UV light, available only to the experimenter, was used to ensure that the correct stimuli were presented on each trial.

A computer program prepared a pseudo-random set of trial pairings. On each experimental trial, the experimenter simultaneously presented two stimuli to the subject. Subjects used a keypad to indicate if the stimuli were geometrically the 'same' or 'different'. There was an equal probability of receiving a 'same' pairing as a 'different' during a given trial. There was no time limit placed on the comparison and subjects were instructed to work with the emphasis on precision rather than speed. Subjects were free to manipulate the stimuli in any way with either hand, serially or simultaneously.

Pilot experiments showed that each trial was somewhat lengthy – ca. 7–15 s. This would make it difficult to traverse all of the possible 25×25 factorial comparisons, even once, in a reasonable amount of time. Therefore, we chose a band of nearby frequencies for comparisons. On each 'different' trial, the stimuli differed by a maximum of ± 4 frequency steps – equivalent to ± 0.5 cycles/object. Each subject performed two sessions of 200 trials for a total of 2,400 comparisons across all subjects.

2.1.3. Participants

There were a total of six subjects, all undergraduate students at Skidmore College who participated to fulfill research credit requirements for an introductory psychology class. None had seen the objects before and all were naïve with regard to the purposes of the experiment. Finally, all were free of any neurological or physical problems that would interfere with their haptic exploration of the stimuli. Four of the subjects were male.

2.2. Results and discussion

Fig. 2 presents the results of Experiment 1.

Our first analysis computes the d' for each object relative to all other objects, pooled across subjects. d's range from good in the lower frequency objects (maximum $d' = 3.1$) to more mediocre in the higher frequency conditions (minimum $d' = 1.2$). The mean d' across all stimuli is $\bar{x} = 2.2$ with a standard error of $s_{\bar{x}} = 0.18$, indicating that, as a whole, the stimuli are discriminable. Furthermore, a clear inverse relationship is seen between the complexity of the stimuli and their d' with $r^2 \approx 0.6$. Broadly speaking, as the frequency of the stimuli increases, the ability to discriminate it from nearby stimuli decreases.

To illustrate the relative discriminability between two given stimuli a second analysis based on the relative confusability of two given objects was performed. The right side of Fig. 2 shows the result of this analysis. Each location of the matrix indicates the frequency that two given stimuli were confused with each other – e.g. the number of occurrences where the two 'different' stimuli elicited 'same' responses. Darker locations indicate more

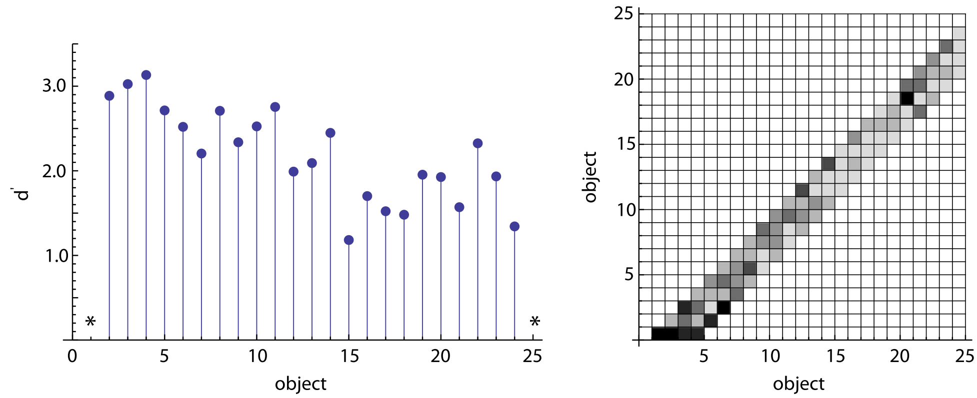

Fig. 2. Results from the unimodal haptic comparison condition. The graph on the left side shows the discriminability of each stimulus. Stars indicate near-infinite (in the case of stimulus 1) and near-zero (stimulus 25) d's. There is a clear inverse relationship between d' and frequency. The right-hand graph shows the relative confusability between two stimuli. Darker squares indicate more frequent confusion between the two indicated stimuli.

confusion. Since the two stimuli were presented simultaneously in the same interval only the lower diagonal matrix is shown. For example, the comparison between stimuli 7 and 8 is the same as between 8 and 7.

A multidimensional scaling on the confusion matrix resulted in a systematic ordering by object frequency, as expected, and a *k*-means cluster analysis found two clusters – a compact group of lower frequency objects from 1 to 5 and a second group consisting of the remaining stimuli, 6 to 25. We suggest that these groupings make phenomenally intuitive sense. Smooth objects do not provide much in the way of 'landmarks' or features that can be used to diagnose relative differences whereas the high frequency objects possess too much information. Confusion within these categories is more frequent than between them.

As in the pilot experiment, these trials took approximately 9 s each for most subjects. However, there was one notoriously meticulous subject who took more than 30 s per trial. Despite this, his results were typical of that of the other subjects.

3. Experiment 2

We proceed by performing a unimodal experiment similar to Experiment 1 using vision instead of touch. Using the same set of objects we determine visual discriminability and confusion.

3.1. Method

The methodology of Experiment 2 is similar to that of Experiment 1 in that subjects perform a discrimination using a single interval two-alternative forced choice (2AFC) paradigm. Two stimuli were presented visually and the subjects were instructed to indicate if the presented objects' shapes were judged to be the same or different.

3.1.1. Apparatus

The apparatus consisted of an Apple Power Mac computer with an ATI 1900-series graphics card, driving a color and luminance-calibrated Apple 30" Pro LCD display. The software was written using an in-house, Python-based, experiment package – *eelpy*. Responses were recorded using an X-keys programmable USB button–box (PI Engineering).

3.1.2. Stimuli

The stimuli consisted of the same 25 objects used in Experiment 1.

A pilot experiment was conducted comparing discrimination performance when using the actual sculpture stimuli versus computer-generated versions. For the stimuli used here, performance was equivalent, therefore we chose to use the computer-generated method since it is significantly more efficient and less time consuming than having an experimenter manually manipulate the stimuli.

Each object was densely sampled at ca. 20,000 triangles/object, painted a uniform, saturated blue and rendered using OpenGL. Simulated illumination was provided by a single source located at the front, upper right quadrant with respect to the viewer. Traditional flat shading with specular highlights was employed with the surfaces' shading coefficients $K_a = 0.01$, $K_d = 0.75, K_s = 0.24$ for the ambient, diffuse, and specular components, respectively.

3.1.3. Procedure

Subjects were seated at a viewing distance of 57 cm (1° = 1 cm) from the LCD monitor. The stimuli were scaled such that their projection onto the screen was consistent with that of a 55 cm^3 object at that viewing distance. In order to enhance the perceived depth of the presented stimuli an eye patch was used to ensure monocular viewing. Subjects were free to pick their preferred eye. Finally, no chinrest was used, allowing for somewhat free viewing, but subjects were instructed to remain relatively still throughout the experiment.

On each experimental trial, two stimuli were presented simultaneously on the monitor, each rotating at a speed of 60° ± 20°/s (Gaussian distributed) about an arbitrary axis constrained to pass through the objects' center of gravity. Subjects used a keypad to indicate if the stimuli were geometrically the 'same' or 'different'. There was no time limit placed on the comparison and subjects were instructed to work with the emphasis on precision rather than speed. Since each stimulus was rotated about an arbitrary axis and had a different rotation rate subjects could not simply compare the stimulus' image information to make their judgements. Finally, on each trial there was an equal probability of receiving a 'same' pairing as a 'different'.

Since these visual comparisons are much faster than the haptic condition we tested a wider range of stimulus pairings. For consistency with our other experiments, our principal analysis is restricted to the same maximum difference of ± 4 frequency steps as in the haptic condition.

Each subject performed two sessions of 500 trials for a total of 13,000 comparisons across all subjects. Due to the large number of possible pairings and a desire to cover the entire range with multiple trials per pairing, a random subset of pairings was chosen for each subject at each session. This method resulted in about 22 judgements per possible pairing condition.

3.1.4. Participants

Experiment 2 used a total of thirteen subjects. The majority (9 subjects) were undergraduate students at Skidmore College who participated to fulfill research credit requirements for an introductory psychology class. The remainder consisted of the authors and colleagues. With the exception of the authors, all had never seen the objects before and were naïve with regard to the purposes of the experiment. All of the subjects possessed normal, or corrected-to-normal, visual acuity. Eight of the subjects were male.

3.2. Results and discussion

Fig. 3 presents the results of Experiment 2. As with Experiment 1 we perform two analyses – one on the individual stimulus discriminability and the other on their relative confusability.

As with Experiment 1, our first analysis computes the d' for each object, pooled across subjects. d' s for the entire set of stimuli range from $d' = 1.0$ to 2.5 with a mean d' of $\bar{x} = 1.8$ and standard error of $s_{\bar{x}} = 0.06$. Unlike in Experiment 1, the pattern of d' as a function of frequency shows a poor linear relationship ($r^2 \approx 0.2$). The overall d' s are significantly lower than the d' s in the haptic condition, suggesting that the visual discriminations were more difficult overall but did not suffer from the frequency effect seen in the haptic case.

The right side of Fig. 3 shows the result of the confusion analysis. Each location of the matrix indicates the frequency that two given stimuli were confused with each other – e.g. the number of occurrences where the two 'different' stimuli elicited 'same' responses. Darker locations indicate more confusion.

A multidimensional scaling of the confusion matrix was performed, yielding a configuration similar to that seen in Experiment 1 – an ordering of the stimuli by frequency. A *k*-means cluster analysis yielded three clusters of confusability – A large group between stimuli 1 and 10, a second large group from about 18 to 25 and a third, smaller group from about 12 to 16. Within these groups there are some stimulus pairs that are clearly more frequently confused. For example, 5 & 6, 20 & 21, and 23 & 25 are often mistaken for each other.

Finally, these trials were completed much more quickly than in Experiment 1, each taking ca. 2–4 s as opposed to the haptic trials' 9 s.

4. Experiment 3

This experiment combines Experiments 1 and 2 into a crossmodal (visual-haptic) task.

4.1. Method

The design is the same simple discrimination using the single interval two-alternative forced choice (2AFC) paradigm used in Experiments 1 and 2. In this experiment the two objects were presented simultaneously, first haptically and a second presented visually. Subjects were instructed to indicate if the presented shapes were judged to be geometrically the same or different.

4.1.1. Stimuli and apparatus

The experiment setup was similar to that of Experiment 2 with the addition of the computer monitor set up as in Experiment 1.

The stimuli consisted of the same 25 objects used in Experiments 1 and 2.

4.1.2. Procedure

As in Experiment 2, a computer program prepared a pseudorandom set of trial pairings. On each experimental trial, one object was presented for haptic exploration and a second was presented visually. Subjects used a keypad to indicate if the objects were geometrically the 'same' or 'different'. There was an equal probability of receiving a 'same' pairing as a 'different' during a given trial. There was no time limit placed on the comparison and subjects were instructed to work with the emphasis on precision rather than speed. Subjects were free to manipulate the haptic stimuli in any way with either or both hands.

Keeping with the constraints used in our prior experiments, we reduced both the number of comparisons per experimental block and reduced the range of frequencies used when comparing. On each 'different' trial, the stimuli differed by a maximum of ±4 frequency steps. Each subject performed two sessions of 200 trials for a total of 2,400 comparisons across all subjects.

4.1.3. Participants

There were a total of six subjects, all undergraduate students at Skidmore College who participated to fulfill research credit requirements for an introductory psychology class. None had served in previous experiments or had seen the objects before and all were naïve with regard to the purposes of the experiment. Finally, all were free of any neurological or physical problems that

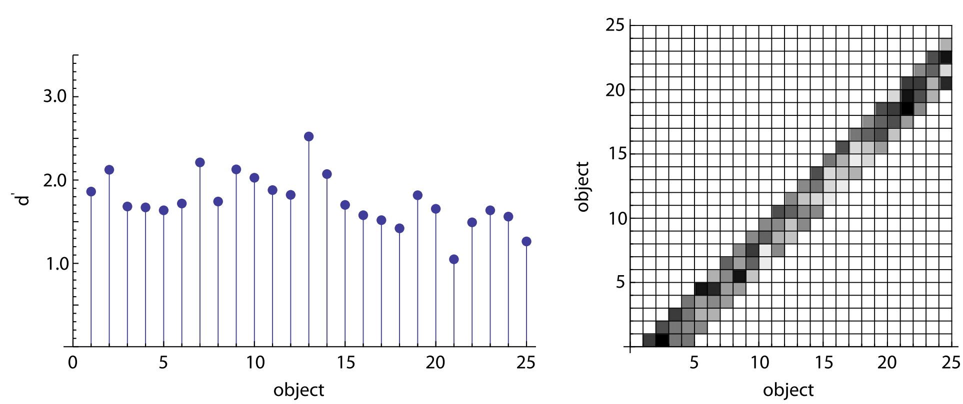

Fig. 3. Results from the unimodal visual comparison condition. The graph on the left side shows the discriminability of each stimulus. Recall that the spatial frequency of surface features increases as the objectnumber increases. There is no meaningful linear relationship between d' and frequency. The right-hand graph shows the relative confusability between two stimuli. Darker squares indicate more frequent confusion between the two indicated objects.

would interfere with their haptic exploration of the stimuli. Four of the subjects were male.

4.2. Results and discussion

As with Experiments 1 and 2 we performed two analyses – one on the individual stimulus discriminability and the other on their relative confusability.

Results from the crossmodal visual-haptic comparison condition are shown in Fig. 4. The graph on the left side shows the discriminability of each object. Stars indicate near-zero d's. As with the haptic-only condition, there is a linear relationship between d' and frequency ($r^2 \approx 0.5$). However, there seems to be a bimodality in the results with d's dropping off precipitously with the higher frequency objects, starting with object 17 and continuing to the end of the stimulus range. This condition had a mean d', $\bar{x} = 1.79$ with $s_{\bar{x}} = 0.01$ – similar to the results found in the unimodal haptic condition.

The right-hand graph shows the relative confusability between two objects. Darker squares indicate more frequent confusion between the two indicated objects. Here we see a rather distributed confusion compared to the unimodal conditions. This is reinforced by the findings of a *k*-means cluster analysis that shows no significant grouping of stimuli.

Trials were completed at roughly the same rate as the unimodal haptic condition – approximately 9 s/trial.

These results suggest that the information is able to be integrated across sensory modalities and therefore perceptual equivalence exists. However, this is not without significant error as frequency increases.

5. Discussion

The previous experiments undertook the task of investigating our ability to discriminate objects based on haptic and/or visual perceptions. Perceptual equivalency certainly exists between the visual and haptic perceptual modalities, but it is by no means exact.

Overall performance between the three experiments is surprisingly similar. Fig. 5 shows a summary of the results across Experiments 1–3. Prior work suggests that performance is stimulus and modality dependent – Haptic discriminations should be better for stimuli with locally defined features, while vision is more finely tuned for stimuli with global features. Our results show that performance is roughly the same for the unimodal visual and crossmodal conditions while the haptic condition had slightly better performance. This result is the opposite of what would be predicted based on our original theoretical definition of the stimuli as being globally unique. A better description of our stimulus set may be that they are within a range located in the middle of the local-global spectrum. This is not a far-fetched claim. Our stimuli are not as globally unique as bell peppers, nor could their features be considered locally defined in a way similar to textures, but rather somewhere in between.

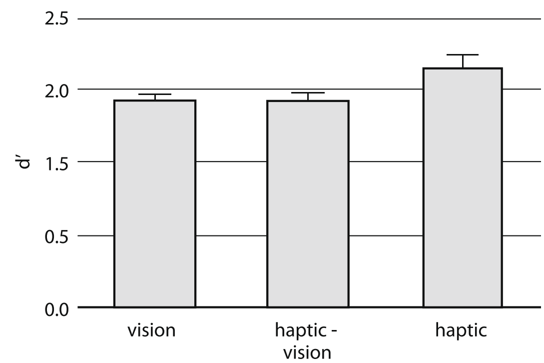

Fig. 5. Results from Experiments 1 to 3. The d' for the three conditions shows that performance is roughly equivalent for the visual discrimination task and the crossmodal discrimination tasks. Haptic discriminations were significantly easier. (Error bars indicate one $s_{\bar{x}}$.)

By using stimuli whose complexity, as defined by the scale and frequency of features, varied in a systematic manner, we could further investigate the effect of complexity on performance. We found that visual discrimination remained mostly unaffected by changes in object complexity ($r^2 \approx 0.2$). However, as the complexity increased, the tasks became precipitously more difficult for the haptic discrimination task ($r^2 \approx 0.6$) and to less of an extent for the crossmodal discrimination task ($r^2 \approx 0.5$).

Fig. 6 shows the linear regression model for each of the three experiments. For stimuli on the lower end of the spatial frequency spectrum, the haptic discrimination task has the best performance while visual performance suffers comparatively. The stimuli in this lower range may have diagnostic features that are only apparent to the haptic system. Crossmodal performance falls roughly halfway between the two unimodal tasks. If it were the case that perceptual equivalence is limited by the weakest link, we would expect crossmodal discrimination performance to be only as good as in the vi-

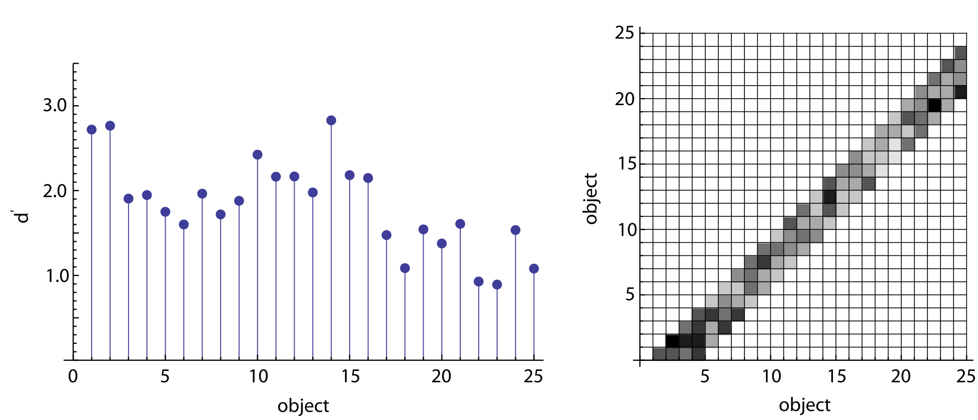

Fig. 4. Results from the crossmodal visual-haptic comparison condition. The graph on the left side shows the discriminability of each object. As with the haptic-only condition, there is a relationship between d' and frequency. There further appears to be a bimodality in the results with d' dropping off precipitously with the higher frequency stimuli. The right-hand graph shows the relative confusability between two objects. Darker squares indicate more frequent confusion between the two indicated objects. Here we see a more distributed confusion compared to the unimodal conditions.

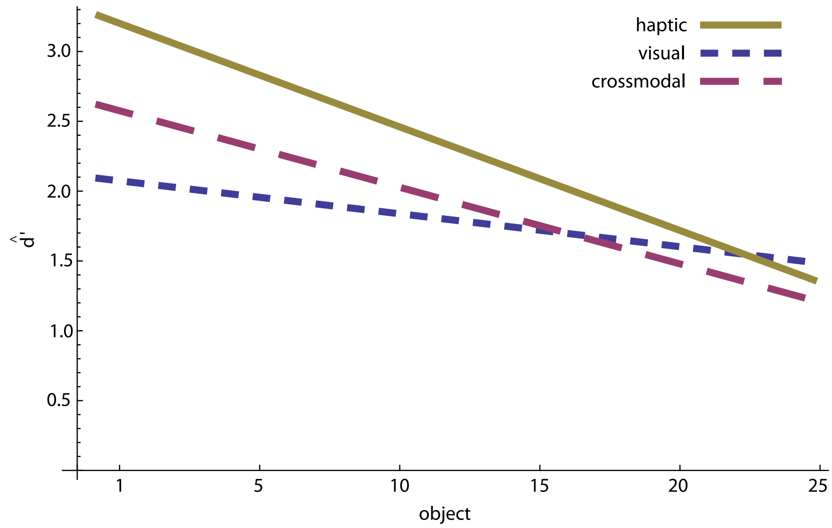

Fig. 6. Fitted models from all three conditions for the relationship between discrimination performance and spatial frequency. The haptic and crossmodal fits had relatively large r^2s, ≈ 0.6 and ≈ 0.5, respectively. The visual condition had a much lower $r^2 \approx 0.2$.

sual condition. One explanation of this result is that the haptic system is bringing the visual systems focus to diagnostic features that the visual system may have otherwise overlooked. Norman et al. (2004) suggested that haptic perception may be more sensitive to local, phase-independent information. In those experiments, subjects were often able to correctly haptically discriminate between two bell peppers when diagnostic local features were present. For example, subjects often credited their correct identifications to specific fold or stem that made a bell pepper unique. Our stimuli were largely devoid of such landmark features, but our results do not rule out this possibility either. As the number of features increases, it would become more difficult to find and then correctly identify or rule out the existence of a haptically located diagnostic feature.

Recent research by Wijntjes, van Lienen, Verstijnen, and Kappers (2008), Wijntjes, Volcic, Pont, Koenderink, and Kappers (2009) claims that "touch disambiguates vision". It is well established that the perceptions of the visual system are often ambiguous within a scaling in depth (Todd, 2004). This may lead us to base second order shape information judgements solely on haptic perceptions when available (Wijntjes et al., 2008; Wijntjes et al., 2009). Similarly, it has been proposed that first order shape information, such as orientation, is integrated in a statistically optimal manner (Ernst & Bülthoff, 2004). But others argue this may not be the case (Lacey et al., 2007; Rosas, Wagemans, Ernst, & Wichmann, 2005; Rosas, Wichmann, & Wagemans, 2007). The exact nature of sensory integration is not critical to the current research, which deals more with perceptual equivalency. According to the current results, the extra perceptual information that arises from utilizing two perceptual modalities increases crossmodal performance. This is only the case for a subset of our stimuli with relatively low spatial frequencies. Further research will be needed to further expand our knowledge of perceptual equivalency within the larger domains 3D shape.

The nature of the task is another possible explanation to why crossmodal performance is higher than the unimodal vision performance. In the crossmodal task each perceptual system only needs to attend to a single stimuli at a time. That is true only for the crossmodal experiment, and a slight downside to a simultaneous discrimination task. It might be expected that this advantage would continue for the entire range of stimuli. As you can see, performance differences disappear as stimuli become more complex. At the higher spatial frequency range, performance appears to be equal for all conditions.

The dependance of perceptual equivalence on complexity is further supported by our clustering analyses. The stimuli appear to be clustered into two or three groups based on the modality employed. Vision yielded three groups (low, mid, and high frequency) whereas the haptic condition yielded two (low and mid-high). In the crossmodal condition the results were much noisier. If the two modalities were completely equivalent, we would expect the clusters to be the same. According to Garbin and Bernstein (1984) stimuli which are perceptually grouped together will be more perceptually equivalent. Although we did not perform an explicit perceptual grouping task in the current experiments, our cluster analyses data support this theory. A low frequency cluster was found for both unimodal conditions. As expected crossmodal performance was better for stimuli with lower spatial frequencies. We therefore conclude that perceptual equivalency is dependent on stimulus spatial frequency for the range of 3D shapes investigated.

To return to Molyneux' question, he was correct to assume a high degree of perceptual equivalence between vision and haptics. It should also be noted that when attempting to teach the newly sighted to identify objects, it may be best to start with simple objects.

References

Amedi, A., Malach, R., Hendler, T., Peled, S., & Zohary, E. (2001). Visuo-haptic object-related activation in the ventral visual pathway. *Nature Neuroscience, 4*(3), 324–330.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94*(2), 115–147.

Blake, R., Sobel, K., & James, T. (2004). Neural synergy between kinetic vision and touch. *Psychological Science, 15*(6), 397–402.

Bruce, C., Desimone, R., & Gross, C. (1981). Visual properties of neurons in a polysensory area in superior temporal sulcus of the macaque. *Journal of Neurophysiology, 46*(2), 369–384.

Caviness, J. A. (1962). *The equivalence of visual and tactual stimulation for solid shape perception*. Master's Thesis. Cornell.

Caviness, J. A. (1964). *Visual and tactual perception of solid shape*. Ph.D. Thesis. Cornell.

De Winter, J., & Wagemans, J. (2006). Segmentation of object outlines into parts: A large-scale integrative study. *Cognition, 99*(3), 275–325.

Duhamel, J., Colby, C., & Goldberg, M. (1998). Ventral intraparietal area of the macaque: Congruent visual and somatic response properties. *Journal of Neurophysiology, 79*, 126–136.

Ernst, M., & Bülthoff, H. (2004). Merging the senses into a robust percept. *Trends in Cognitive Science, 8*(4), 162–169.
Feinberg, T. (1986). Multimodal agnosia after unilateral left hemisphere lesion. *Neurology, 36*(6), 864.
Garbin, C. (1988). Visual-haptic perceptual nonequivalence for shape information and its impact upon cross-modal performance. *Journal of Experimental Psychology: Human Perception and Performance, 14*(44), 547–553.
Garbin, C. (1990). Visual-touch perceptual equivalence for shape information in children and adults. *Perception & Psychophysics, 48*(3), 271–279.
Garbin, C. P., & Bernstein, I. H. (1984). Visual and haptic perception of three-dimensional solid forms. *Perception & Psychophysics, 36*(2), 104–110.
Gibson, J. J. (1962). Observations on active touch. *Psychological Review, 69*, 477–491.
Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston: Houghton Mifflin.
Heller, M., & Clyburn, S. (1993). Global versus local processing in haptic perception of form. *Bulletin of the Psychonomic Society, 31*(6), 574–576.
Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition, 18*, 65–96.
Hoffman, D. D., & Singh, M. (1997). Salience of visual parts. *Cognition, 63*, 29–78.
James, T. W. (2003). Ventral occipital lesions impair object recognition but not object-directed grasping: An fMRI study. *Brain, 126*(11), 2463–2475.
Kappers, A. M. L., Koenderink, J. J., & Oudenaarden, G. (1997). Large scale differences between haptic and visual judgments of curvature. *Perception, 26*(3), 313–320.
Kayaert, G., & Wagemans, J. (2009). Delayed shape matching benefits from simplicity and symmetry. *Vision Research, 49*(7), 708–717.
Lacey, S., Campbell, C., & Sathian, K. (2007). Vision and touch: Multiple or multisensory representations of objects? *Perception, 36*(10), 1513–1521.
Lakatos, S., & Marks, L. (1999). Haptic form perception: Relative salience of local and global features. *Perception & Psychophysics, 61*(5), 895–908.
Locke, J. (1694). *An essay concerning human understanding*. Oxford: Clarendon Press. 1979 ed..
Locke, J. (1978). *The correspondence of John Locke*. Oxford: Clarendon Press.
Macaluso, E., & Driver, J. (2001). Spatial attention and crossmodal interactions between vision and touch. *Neuropsychologia, 39*, 1304–1316.
Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology, 9*, 353–383.
Norman, J. F., Bartholomew, A. N., & Burton, C. L. (2008). Aging preserves the ability to perceive 3d object shape from static but not deforming boundary contours. *Acta Psychologica, 129*(1), 198–207.
Norman, J. F., Clayton, A. M., Norman, H. F., & Crabtree, C. E. (2008). Learning to perceive differences in solid shape through vision and touch. *Perception, 37*(2), 185–196.
Norman, J. F., Crabtree, C. E., Norman, H. F., Moncrief, B. K., Herrmann, M., & Kapley, N. (2006). Aging and the visual, haptic, and cross-modal perception of natural object shape. *Perception, 35*(10), 1383–1395.
Norman, J. F., Norman, H. F., Clayton, A. M., Lianekhammy, J., & Zielke, G. (2004). The visual and haptic perception of natural object shape. *Perception & Psychophysics, 66*(2), 342–351.
Norman, J. F., Phillips, F., & Ross, H. E. (2001). Information concentration along the boundary contours of naturally shaped solid objects. *Perception, 30*(11), 1285–1294.
Op de Beeck, H. P., Torfs, K., & Wagemans, J. (2008). Perceived shape similarity among unfamiliar objects and the organization of the human object vision pathway. *Journal of Neuroscience, 28*(40), 10111–10123.
Ostrovsky, Y., Andalman, A., & Sinha, P. (2006). Vision following extended congenital blindness. *Psychological Science, 17*(12), 1009–1014.
Phillips, F. (2004). Creating noisy stimuli. *Perception, 33*(7), 837–854.
Phillips, F., Thompson, C. H., & Voshell, M. G. (2004). Effects of 3-D complexity on the perception of 2-D depictions of objects. *Perception, 33*(1), 21–33.
Phillips, F., Todd, J. T., Koenderink, J. J., & Kappers, A. M. L. (2003). Perceptual representation of visible surfaces. *Perception & Psychophysics, 65*(5), 747–762.
Picard, D. (2006). Partial perceptual equivalence between vision and touch for texture information. *Acta Psychologica, 121*(3), 227–248.
Rock, I., & Harris, C. S. (1967). Vision and touch. *Scientific American, 216*(5), 96–104.
Rock, I., & Victor, J. (1964). Vision and touch: An experimentally created conflict between the two senses. *Science, 143*(3606), 594–596.
Rosas, P., Wagemans, J., Ernst, M., & Wichmann, F. (2005). Texture and haptic cues in slant discrimination: Reliability-based cue weighting without statistically optimal cue combination. *Journal of the Optical Society of America A, 22*(5), 801–809.
Rosas, P., Wichmann, F., & Wagemans, J. (2007). Texture and object motion in slant discrimination: Failure of reliability-based weighting of cues may be evidence for strong fusion. *Journal of Vision, 7*(6), 1–21.
Sathian, K., Zangaladze, A., Hoffman, J., & Grafton, S. (1997). Feeling with the mind's eye. *NeuroReport, 8*(18), 3877–3881.
Tiest, W. M. B., & Kappers, A. M. L. (2006). Analysis of haptic perception of materials by multidimensional scaling and physical measurements of roughness and compressibility. *Acta Psychologica, 121*(1), 1–20.
Tiest, W. B., & Kappers, A. (2007). Haptic and visual perception of roughness. *Acta Psychologica, 124*(2), 177–189.
Todd, J. T. (2004). The visual perception of 3d shape. *Trends in Cognitive Science, 8*(3), 115–121.
Vogels, I. M., Kappers, A. M. L., & Koenderink, J. J. (1997). Investigation into the origin of the haptic aftereffect of curved surfaces. *Perception, 26*(1), 101–117.
Wijntjes, M. W. A., van Lienen, T., Verstijnen, I. M., & Kappers, A. M. L. (2008). Look what i have felt: Unidentified haptic line drawings are identified after sketching. *Acta Psychologica, 128*(2), 255–263.
Wijntjes, M. W. A., Volcic, R., Pont, S. C., Koenderink, J. J., & Kappers, A. M. L. (2009). Haptic perception disambiguates visual perception of 3d shape. *Experimental Brain Research, 193*(4), 639–644.
Zangaladze, A., Epstein, C., Grafton, S., & Sathian, K. (1999). Involvement of visual cortex in tactile discrimination of orientation. *Nature, 401*, 587–590.