Perception, 2008, volume 37, pages 1867 – 1878

doi:10.1068/p6038

Canonical views in haptic object perception

Andrew T Woods§, Allison Moore, Fiona N Newell¶
School of Psychology and Institute of Neuroscience, Lloyd Building, Trinity College Dublin, Dublin 2, Ireland; e-mail: fiona.newell@tcd.ie; § now at Unilever R&D, The Netherlands
Received 12 February 2008, in revised form 20 June 2008; published online 27 November 2008

Abstract. Previous investigations of visual object recognition have found that some views of both familiar and unfamiliar objects promote more efficient recognition performance than other views. These views are considered as canonical and are often the views that present the most information about an object's 3-D structure and features in the image. Although objects can also be efficiently recognised with touch alone, little is known whether some views promote more efficient recognition than others. This may seem unlikely, given that the object structure and features are readily available to the hand during object exploration. We conducted two experiments to investigate whether canonical views existed in haptic object recognition. In the first, participants were required to position each object in a way that would present the best view for learning the object with touch alone. We found a large degree of consistency of viewpoint position across participants for both familiar and unfamiliar objects. In a second experiment, we found that these consistent, or canonical, views promoted better haptic recognition performance than other random views of the objects. Interestingly, these haptic canonical views were not necessarily the same as the canonical views normally found in visual perception. Nevertheless, our findings provide support for the idea that both the visual and the tactile systems are functionally equivalent in terms of how objects are represented in memory and subsequently recognised.

1 Introduction

Previous research on visual object recognition has found evidence for a preference in encoding and representing objects from specific but consistent viewpoints, referred to as canonical views of objects (Blanz et al 1999; Palmer et al 1981; Perrett et al 1992). In turn, these viewpoints promote better subsequent recognition of the objects: recognition performance is typically more accurate and less time-consuming than recognising objects from other views (eg Cutzu and Edelman 1994; Newell et al 2001; Verfaillie and Boutsen 1995). Here we investigated whether preferences for certain positions of objects also exist for objects encoded through touch only and, if so, whether the objects are more easily recognised from these object positions relative to other object positions.

Many have argued that in visual object recognition objects are encoded and stored in memory as view-specific representations. Indeed, much evidence has been provided in support of a view-specific representation of objects, from animal physiology (Logothetis and Pauls 1995), to human behavioural (eg Hayward and Williams 2000; Newell and Findlay 1997; Tarr and Bülthoff 1998), and neuroimaging studies (eg Gauthier et al 2002; see also Peissig and Tarr 2007 for a review). Accordingly, visual recognition relies on the successful matching of an encoded view of an object to either the stored view or one of a number of stored views of that object in memory, and recognition is most efficient when the encoded view is very similar to these stored representations (eg Edelman and Bülthoff 1992; Tarr and Bülthoff 1998).

Canonical views in vision are thought to arise either as a consequence of view-dependent representations or because they allow efficient access to the structural description of objects in memory. For example, these views may provide access to an optimal amount of information about object structure (Biederman 2000; Biederman and

¶ Author to whom all correspondence should be addressed.

Gerhardstein 1993) or object features (Ullman 1998) or may include the most optimal amount of information in its 2-D view to allow for greater generalisability across novel views (Edelman and Weinshall 1991). Canonical views can be interpreted as views of an object that are most likely to match the stored representations of objects in visual memory and subsequently facilitate more efficient object recognition. Moreover, canonical views seem to be independent of object familiarity as they have been demonstrated for both familiar (eg Palmer et al 1981) and unfamiliar (eg Cutzu and Edelman 1994) objects. What is not clear, however, is whether canonical, or preferred, views of objects are also found in haptic object recognition. Given the differences between vision and touch in the encoding and, possibly, representation of object information, we were interested in determining whether the principle of canonical views in visual object recognition also holds true for objects encoded through touch, ie whether there are 'views' of objects that are consistently preferred in haptics and whether these views specifically promote more robust haptic object recognition.(1)

There is behavioural evidence to suggest that vision and haptics share common principles of information encoding and storage for the purpose of object perception (Easton et al 1997; Norman et al 2004; Reales and Ballesteros 1999) and, as such, canonical views may be a characteristic of both modalities. Neuroimaging evidence, moreover, suggests a degree of overlap between the cortical structures involved in both visual and haptic object perception (Amedi et al 2001; James et al 2002). For example, both Amedi et al and James et al find evidence for activation for haptic object recognition in the lateral occipital complex, an area known to be involved in visual object recognition (eg Malach et al 1995). Together these findings suggest that there is a great deal of overlap between the functional manner in which object information is processed and also the neural processes underlying object perception through vision and touch.

Although vision and touch may process object information in a similar manner for the purpose of object representations, vision and haptics encode object information in very different ways and the differences in encoding may affect crossmodal differences in the nature of the information that is stored. For example, visual encoded views are restricted to the information contained in the 2-D image projected onto the retina of the object facing the observer. As such, images containing the optimal amount of information about the object's structure and feature composition may be preferred and will likely promote more efficient recognition. In contrast, haptic object encoding is not restricted to a 2-D projected image; rather, all object information required for recognition purposes can be encoded rapidly, at a 'glance', seemingly irrespective of object orientation (eg Klatzky and Lederman 1995). If this is so, then object orientation may be unimportant for haptic object recognition and canonical views of objects encoded through touch may not arise. However, it has been shown that if the position of the object is constrained then preferred orientations may arise as some hand postures are more comfortable to maintain than others and 'awkward' postures can lead to inefficient encoding of the object stimulus (eg Heller et al 2003; Kappers and Viergever 2006). Furthermore, such encoding preferences can lead to more efficient haptic recognition of certain 'views' of novel objects over others (Newell et al 2001). Recently, Ernst et al (2007) reported that the free exploration of spatially unanchored objects by means of touch is orientation specific: that is, participants consistently constrained the orientation of a novel object during haptic exploration even in the absence of any explicit instructions on how to explore the objects. As such, evidence is emerging to suggest that objects are encoded through touch on the basis of a limited set of views,

(1) By object 'view' we mean the position of the object relative to the participant. Although the term object 'view' is not appropriate to haptics, we are using it here simply to maintain consistency of terminology with the literature on visual object recognition.

irrespective of whether the object is spatially anchored or not. What is not yet known is whether there are preferred views of objects that are consistent across participants.

In our study here we investigated two specific questions: first, whether or not there was evidence for preferred orientations of both familiar and novel objects encoded through touch that were consistent across participants (experiment 1); and, second, whether these views afforded more efficient object recognition relative to other non-preferred views (experiment 2). We were also curious to determine whether canonical views in haptic object recognition bore any similarity to the ¾ canonical view often reported in visual object perception (eg Palmer et al 1981).

2 Experiment 1

Our aim here was to identify views of objects in haptic object perception that were preferred by each participant and to determine if these views were consistently chosen across individuals. If there are preferred views for haptic object perception, then we expected very little variation in terms of the positions in which the participants place the objects (as found by Palmer et al 1981 for the visual encoding of objects). On the other hand, if objects are nevertheless encoded through specific, preferred views for each participant, these views would not be considered canonical if there is little consistency (ie high variability) in preferred viewpoints of objects across participants.

In order to reduce the possibility that visual knowledge (eg imagery) plays a role in the choice of preferred views in haptic perception, we tested the consistency of preferred views for both familiar and novel objects. If vision gives rise to canonical views in haptic object perception, then we were likely to find consistency across views for familiar but not for unfamiliar objects (since these latter objects are less easily imagined than familiar objects).

2.1 *Method*

2.1.1 *Participants.* Twelve undergraduate and postgraduate students (seven male and five female) from Trinity College Dublin volunteered to participate in this experiment in return for either course credit or payment. Participants' ages ranged from 23 to 55 years with a mean age of 29 years. Two of the participants were left-handed and none reported any tactile impairment. All participants provided written informed consent prior to the experiment and our protocol was approved by the School of Psychology Research Ethics Committee.

2.1.2 *Stimuli and apparatus.* Ten familiar objects and six novel objects were used as stimuli. The familiar objects comprised a set of common objects such as children's toys, and realistic small-scale models of real objects. These objects were a model horse, car, teapot, house, bicycle, telephone, iron, rocking chair, hairdryer, and a shoe. We created the novel objects using DUPLO™ bricks and the objects were based on a design of unfamiliar objects used in previous research (eg Ernst et al 2007; Newell et al 2001). Each novel object was constructed with six DUPLO™ bricks and each was based on a unique configuration of these bricks. In order to render the novel objects similar to the familiar objects, all novel objects had one common feature: one side of each object was completely flat (similar, for example, to the underside of a telephone or iron). Apart from that flat surface, we carefully designed our novel objects to avoid a bias in which one surface of the object was more informative of object identity than another. See figure 1 for a depiction of all objects used in our study.

The participant was seated in front of a table on which each object was to be finally positioned. A curtain was used to occlude the objects from the participant's sight at all times. Once the participant placed the object into his or her preferred position, a digital camera (JVC Digital Mini-DV camera) was used to record the image of the object [note that the camera was positioned opposite the participant so that

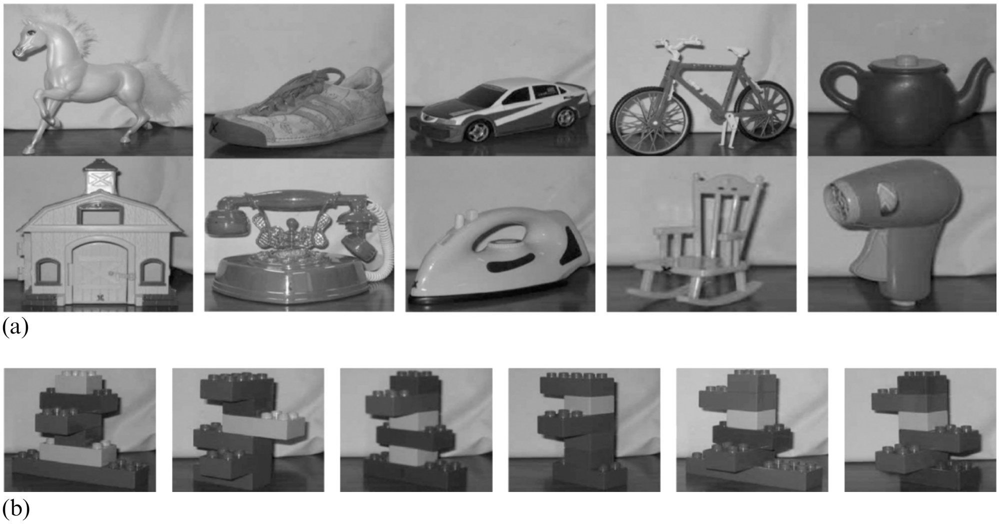

Figure 1. An illustration of familiar (a) and novel (b) objects used in our experiments.

object images were consequently of views facing away from the participant (see Newell et al 2001)]. The camera was held in place with a tripod, so that it faced the object, curtain, and participant.

2.1.3 *Design.* The experiment was based on a within-subject design with object familiarity (familiar or novel) as the main factor. The dependent variable in this study was the 3-D orientation at which participants placed each object. In order to ensure that the novel objects were not influenced by the presence of the familiar objects in the task, the experiment consisted of two blocks of trials with one containing familiar objects only and the other novel objects only. Block order was counterbalanced and trial order was randomised across participants.

2.1.4 *Procedure.* Before we tested whether there are indeed preferred views of objects encoded through haptics, we first had to determine whether our set of familiar objects was recognisable to participants. We therefore conducted a pilot study with ten naive participants (seven females and three males with an average age of 25.9 years, nine of whom were right-hand dominant) on their ability to recognise the ten familiar objects used in the main experiment. With the same apparatus as described above, participants were handed each of the objects one at a time in random sequence, out of sight, and were asked to name each object as fast and as accurately as possible. The object names were recorded by the experimenter who also timed the responses using a stopwatch. A response was coded as correct if it matched the basic name of the object (eg 'bicycle', 'house', and 'hairdryer' were marked as correct answers as were 'bike', 'toy house', and 'dryer for hair'). Examples of incorrect responses provided to the same objects were 'carriage', 'measuring tape', and 'toy gun', respectively. We found that, on average, each participant correctly named 9.2 (SE $= 0.29$) of the familiar objects. When errors were made, they were not consistently made on any one object. Furthermore, it took the participants an average of 10.26 s (SE $= 1.12$ s) to provide the correct name of each object. We were confident that our familiar objects were easily identifiable through touch.

In the main experiment, participants were presented with an object directly into their open hands in each trial which they were instructed to freely explore with both hands and then place on the table in its most optimal position for learning. In order to avoid any bias in the position of the object, each object was presented in a random orientation.

Participants were instructed to position the object on the table in such a manner that, if someone were to learn the object for the first time (eg a visually impaired child), this object position would be the most optimal and informative for haptic learning and subsequent recognition. They were also required to position the object as if it were placed in front of the person learning the object. There was no time limit for this task and no restrictions given on object exploration. There were 16 trials in the experiment (corresponding to 1 trial per object). Participants indicated the end of a trial by removing their hands from the positioned object. Each trial began as soon as an object was placed in front of the participant. An image record of the object position was taken by the experimenter at the end of each trial. The experiment lasted for approximately 15 min.

2.2 *Results*

Using the photographic images of the objects as reference, we first noticed that there was a large amount of consistency across participants with regard to the preferred position of each object. Initially we decided to categorise each of the object orientations in step sizes of 45° along each of the X, Y, and Z axes, yielding 512 total possible discrete views. Once each of the object positions was categorised into one of the 512 possible orientations, these object positions were then ranked in terms of their frequency across participants. For example, the shoe was placed at 5 different orientations across all twelve participants: six placed the shoe so that its tip pointed towards the participant's right; three positioned it so that its tip pointed left; and three participants placed the shoe at other positions within the 512 possible views. (Data for one of the unfamiliar objects could not be included in the combined analysis as, by mistake, the shape of the object differed across some of the participants.)

The exact probability for an object being placed N times in the same orientation by all participants was then calculated on the basis of multinomial theory. For this calculation we had to set the total possible number of different views of the objects. As a conservative measure, we decided not to use 512 views highlighted previously, as objects were mostly positioned in an upright freestanding orientation resulting in orientations in, for example, the image plane (ie Z axis) being largely redundant. Because of this we set the number of discrete views for this calculation to 32 (there were 8 potential views along the y-axis, 2 along the x-axis, and 2 along the z-axis for all objects in our experiments). By way of confirmation of our results, we also conducted a Monte-Carlo-type simulation consisting of 10^7 'virtual' experiments to estimate the probabilities. The estimated and exact probabilities were matched for each object, confirming that the complex multinomial theory probability calculations were accurate.

Figure 2 shows the orientations at which each object was placed most commonly and shows the exact probabilities for the frequency of the orientations occurring by chance (ie according to multinomial theory, the minimum number of participants choosing an object view by chance is three). Note that for all the objects there was at least one orientation in which the object was placed by participants at a frequency rating far exceeding that predicted by chance.

We then calculated whether familiar objects were more likely to be positioned in certain views than novel objects. We found no difference in terms of the frequency in which familiar or novel objects were consistently positioned in a preferred view (Mann–Whitney U-test: $U = 13.5$, $p = 0.14$). This finding suggests that visual experience with the object cannot exclusively account for the consistency across participants in which the object was placed. However, we used a relatively small number of novel objects, therefore it remains possible that differences exist between canonical views for familiar and unfamiliar objects which we failed to find.

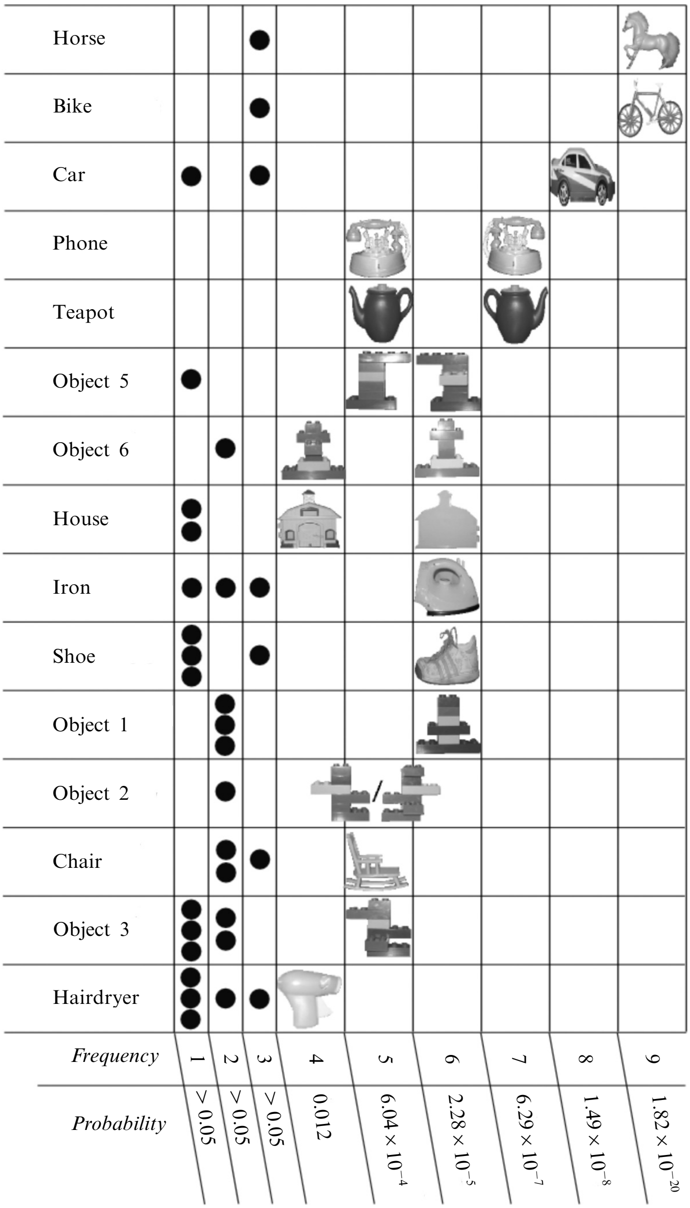

Figure 2. The frequency (ie number of participants) at which the certain orientations for each of our object stimuli were consistently presented. Objects in the figure are presented in rank order from the top, according to the highest number of consistent orientations given per object. Exact probabilities for each of the orientation frequencies are as follows: 4 ($p = 0.0118$); 5 ($p = 0.000604$); 6 ($p = 2.28 \times 10^{-5}$); 7 ($p = 6.29 \times 10^{-7}$); 8 ($p = 1.49 \times 10^{-8}$), and 9 ($p = 1.82 \times 10^{-10}$). Only those orientations of objects which were given at frequencies that exceeded that expected by chance ($p < 0.05$) are displayed in the figure. Some common orientations were chosen of objects but the frequency of these orientations was lower than expected by chance. For clarity, common orientations which did not exceed chance levels are depicted as black dots, and each dot represents a particular object orientation. Some objects had two common orientation frequencies that exceeded chance levels and these object orientations are both included in the plot. One novel object was excluded from the analysis as the result of experimental error and so has not been illustrated.

2.3 *Discussion*

We found that objects were positioned into preferential or canonical views that were largely consistent across participants and at a frequency rate that was greater than purely by chance. Moreover, consistent preferred views were found for both familiar and unfamiliar objects, suggesting that it is unlikely that visual experience was solely responsible for these views.

Interestingly, the canonical views of the objects resulting from haptic object exploration were quite different from those generally identified as visual canonical views (eg Palmer et al 1981; Perrett et al 1992). Visual canonical views of familiar objects are typically presented in a ¾ view (ie 45°) away from the alignment of the object's major axis with respect to the observer. Moreover, visually preferred views of unfamiliar objects also share these ¾-view characteristics of familiar objects (eg Palmer et al 1981). In contrast, we found that the preferred positions of objects explored through active touch (both familiar and unfamiliar objects) were such that each object's elongated axis was aligned either perpendicular or parallel to the participant's body mid-line. Object familiarity, however, did have some role to play in that the preferred views were not completely arbitrary: most familiar objects were positioned in an upright manner (eg the car was placed so that its four wheels touched the ground, the shoe was placed with the opening facing upwards, etc), while unfamiliar objects were typically orientated so that their elongated axis was aligned to the vertical.

We also noted that there was a left – right bias for many of the familiar objects positioned into their canonical views. For example, the horse, bicycle, and shoe were mostly aligned such that the front of the object pointed rightward with regards to the participant, while the car and teapot were orientated leftward. It is possible that this bias is influenced by object function and/or handedness. For example, a right-handed individual would normally mount a bike so that the front of the bike was positioned to the right of the participant. Likewise, a teapot is normally held by the right hand of a right-handed individual so that the spout of the teapot points left, highlighting the role of handedness and object function in haptic canonical views of familiar objects.

3 Experiment 2

In previous studies on canonical views in visual object recognition, objects presented in their canonical views were recognised more efficiently than objects presented from other, non-canonical views (Cutzu and Edelman 1994; Edelman and Bülthoff 1992; Palmer et al 1981; Perrett et al 1992; Verfaillie and Boutsen 1995). We decided to investigate whether the same holds true for the preferred views demonstrated by participants in our previous experiment. To that end, we tested haptic object recognition performance of objects presented in canonical over other non-preferred views in a new group of participants. To ensure our findings would remain independent from the effects of visual experience we used only unfamiliar objects in this experiment.

Preferred views identified in the previous experiment can be termed canonical if they facilitate object recognition. If preferred views do benefit object recognition performance, then this would offer support for the idea of canonical object views in haptics. On the other hand, if recognition accuracy does not differ across views, then it is likely that the preferred views identified in experiment 1 were chosen for reasons unrelated to efficient object encoding and storage in memory (eg aesthetic reasons) and that this would, in turn, imply that canonical views do not always occur in haptic object perception but are perhaps a byproduct of visual object processing.

3.1 *Method*

3.1.1 *Participants.* Twelve naive participants (eight male and four female) from Trinity College Dublin volunteered to take part in this experiment for research credits or payment. Their ages ranged from 23 to 55 years with a mean age of 27 years. One participant reported being left-hand dominant, whereas all others were right-handed. None reported any tactile impairment and none took part in the previous experiment.

3.1.2 *Stimuli and apparatus.* The six novel objects from experiment 1 were used as stimuli along with a further six similar objects which we created for the purposes of this experiment. These six extra objects were necessary in order to limit the number of times the novel objects from the previous experiment were presented, as repeated presentations may have biased the results. The apparatus was also the same as that in experiment 1. The six object stimuli from the previous experiment were always presented in their canonical views during the present experiment. Our canonical views were those identified by the separate group of participants in experiment 1. The newly-created objects were therefore presented in non-canonical views which were chosen to be qualitatively unlike the canonical views of the other six objects. For example, we noticed that the canonical views of the six objects used in experiment 1 often included the upright position, or the supine position. Consequently these views were avoided for the non-canonical views in this experiment. As a minimum, the non-canonical views used here differed by at least 90° in any one axis from the canonical view of any of the six objects used previously. See figure 3 for an example of some of the objects presented in either canonical or non-canonical views.

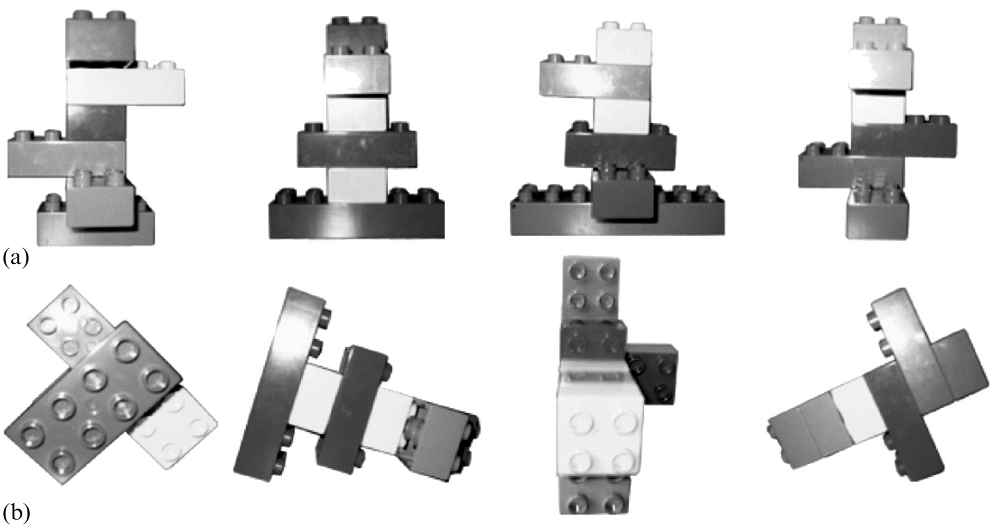

Figure 3. An example of the novel objects used as stimuli in experiment 2. The top row of object images represents objects shown in their canonical views and the lower panel represents the same objects shown in non-canonical views.

3.1.3 *Design.* The experiment was based on a one-factor, within-subjects design with viewpoint as the main factor (canonical or not). The experiment was based on a delayed match-to-sample paradigm and was conducted over two blocks. Matched and mismatched trials were randomly assigned across blocks. Likewise, trials involving matching across canonical or non-canonical views were randomly assigned across blocks. We measured participants' matching performance as the number of correct responses.

3.1.4 *Procedure.* A trial consisted of a sequence of two objects which either matched each other in shape or were different. In each trial, the participant's task was to first learn an object and then to decide whether a second object was the same or different to the previously explored object. In a trial, the first object was placed directly into the participant's hands and he or she was instructed to learn the object by freely exploring it for 30 s. This standard object was presented in a random orientation into the participant's hands. No explicit instructions were given how to explore the object. The second object in a trial was presented immediately after the initial object, but in a fixed position which was either a canonical or non-canonical view. The participant was asked to decide, as fast and as accurately as possible, if the second object was the same as the previous object presented in the trial or not, without moving the object from its position. The participants could not refer to the standard object once the second object was presented. There were 24 trials in total and the experiment took approximately 30 min to complete.

3.2 *Results and discussion*

We calculated the mean number of correct match responses for the object pairs in the 'same' trials (ie hit responses) and the mean number of times they correctly identified the objects as not matching in the 'different' trials (ie correct rejections). For objects that were positioned in a canonical orientation, the mean percentage hits and correct rejections scores were 87.5% and 88.9%, respectively. When objects were positioned in non-canonical orientations, the mean percentage hits and correct rejections scores were 79.1% and 72.2%, respectively. These data were then used to calculate d' scores (a sensitivity measure based on signal detection theory; Green and Swets 1966). The mean d' score across participants for objects which were orientated in a canonical view was 2.14 ($\beta = 1.18$) and the mean d' score for objects oriented in non-canonical views was 1.31 ($\beta = 1.05$). These means are illustrated in figure 4. A pairwise t-test revealed that matching sensitivity to canonically orientated objects was significantly higher than for objects rotated in a non-canonical view ($t_{11} = 3.06$, $p = 0.011$). In order to check for a response bias, we repeated this analysis using β values and found no difference between the conditions ($t_{11} = 0.492$, $p > 0.05$).

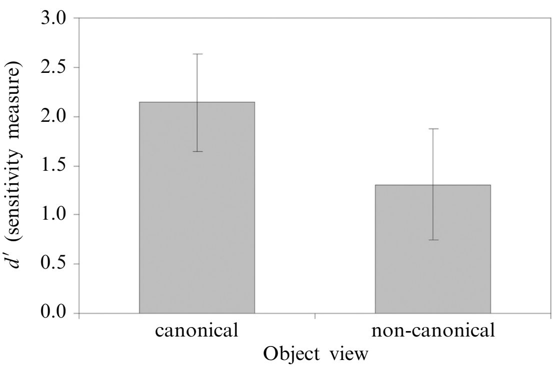

Figure 4. Plot showing the mean d' scores indicating performance across the conditions where the objects were either matched on the basis of their canonical or non-canonical views.

Our data support the idea that preferred views in haptic object perception affect more efficient object recognition than non-canonical views. Specifically, object-matching performance was more accurate for objects presented in canonical views than when objects were orientated in non-canonical positions. Moreover, these canonical views were identified as canonical by a different set of participants (from experiment 1), suggesting the effects of canonical views are robust in haptic (as in visual) object recognition.

4 General discussion

In this study we investigated whether canonical views occur in haptic object perception. We found that there are preferred orientations in which both familiar and novel objects are positioned by means of haptic exploration only and that these orientations were consistent across participants (experiment 1). We also found that these preferred positions or 'views' subsequently promoted better accuracy in object recognition (experiment 2). Our results thus lend support to the idea that canonical views are not restricted to the domain of visual object perception.

Finding evidence for canonical views in haptics provides further support that the processes involved in visual and haptic object recognition are, at least, functionally similar. For example, our data suggest that object information encoded though haptics is based on viewpoint-specific snapshots of the objects. Such information may be the result of specific exploratory procedures used during object encoding (Ernst et al 2007; Lederman and Klatzky 1987). It is possible that haptic encoding of similar objects allows for these viewpoint-specific 'snapshots' to be rapidly encoded at the expense of encoding a full structured, 3-D representation of the object. For objects that are not similar, then, such view-specific encoding would be unnecessary for object recognition since objects may be rapidly discriminable on the basis of a single view or a limited set of encoded features (see, for example, Klatzky and Lederman 1995; Reales and Ballesteros 1999). Encoded views are then compared to the stored representation in memory for the purpose of object recognition (see, for example, Newell et al 2001). One interesting question, raised by the idea of view-specific representations of objects encoded through touch, is whether or not the temporal order at which the object's surfaces are encoded is maintained in memory, as it seems to be for visual object recognition (Wallis and Bülthoff 2001).

Although we found evidence for canonical object views in haptics, our findings also illustrate that qualitative differences exist between visual and haptic canonical views of objects. We found that haptic canonical views are typically orientated so that their main axes are either parallel or perpendicular to the body mid-line. Visual canonical views, on the other hand, are typically a 45° view of the object's main elongated axis relative to the body mid-line (see Palmer et al 1981). Newell et al (2001) similarly found that vision and haptics differ in terms of the orientation at which object surfaces are most accurately encoded: haptic encoding was most proficient for object surfaces that faced away from the observer, whilst visual encoding was best for object surfaces which faced towards the observer. If preferred views of objects for the purposes of object encoding are specific to each modality, then this raises the issue of how object information represented in memory can be efficiently shared across modalities. A recent study suggests that cross-modal object recognition is relatively inefficient, since there is a cost in cross-modal versus within modal recognition performance (Ernst et al 2007). It is possible that this cost may be attributed to differences in the preferred encoded surfaces across vision and touch.

We found no direct evidence that visual experience influenced canonical views in haptic object perception. If visual experience did affect haptic canonical views, then canonical views of familiar objects would likely have been qualitatively different to those of objects that were visually unfamiliar. In other words, if vision affected canonical views through touch, then these views would likely resemble more ¾ views, at least for the more familiar objects. Since no representation exists in visual memory for the novel views, it is unclear whether any consistent preferred views would emerge for these objects. That is not to say that vision had no influence on canonical views through touch, but, if it did, then it is likely to be subtle. For example, we noticed some instances where the canonical view was typically such that the most informative surface of the object faced the observer rather than faced away, which is more conducive to efficient visual, rather than haptic, object recognition (Newell et al 2001). Thus, it is as if in some cases the participants oriented the objects to provide a better view of the most informative surface to the mind's eye.

Unlike our results here, where vision does not seem to influence haptic canonical views, recent literature on the haptic recognition of faces suggests that vision plays a strong role in the canonical views of faces. For example, Kilgour and Lederman (2006) reported a face-inversion effect in haptic face recognition using face masks, which reflects the well-known effect found in the visual recognition of faces (eg Valentine 1988) and may be determined by visual experience. On the other hand, Casey and Newell (2005) reported no difference in haptic recognition performance to face masks presented either facing or turned away from the participant, suggesting that the similarities between visual and haptic recognition of face views may be limited to orientations in the picture plane and not rotations in depth (eg see Bruce et al 1987).

In sum, we found evidence for canonical views in haptic object perception for both familiar and unfamiliar objects. Moreover, these views facilitated better recognition performance relative to other views of objects. In our study we used various shapes of familiar objects but limited our object shapes in the unfamiliar object category. Future research will likely determine whether these canonical views generalise across other categories of familiar and unfamiliar object shapes, such as tetrahedra or 'amoeba'-like shapes. Nevertheless, the data presented here strongly support the idea that, like vision, haptic performance benefits from certain views of objects than others. The implications of these findings stretch beyond basic research and could affect applied research in areas such as the design of haptic devices for virtual-reality and consumer packaging research.

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** 324 – 330

Biederman I, 2000 "Recognizing depth-rotated objects: a review of recent research and theory" *Spatial Vision* **13** 241 – 253

Biederman I, Gerhardstein P C, 1993 "Recognizing depth-rotated objects: evidence and conditions for three-dimensional viewpoint invariance" *Journal of Experimental Psychology: Human Perception and Performance* **19** 1162 – 1182

Blanz V, Tarr M J, Bülthoff H H, 1999 "What object attributes determine canonical views?" *Perception* **28** 575 – 599

Bruce V, Valentine T, Baddeley A D, 1987 "The basis of the 3/4 view advantage in face recognition" *Applied Cognitive Psychology* **1** 109 – 120

Casey S J, Newell F N, 2005 "The role of long-term and short-term familiarity in visual and haptic face recognition" *Experimental Brain Research* **166** 583 – 591

Cutzu F, Edelman S, 1994 "Canonical views in object representation and recognition" *Vision Research* **34** 3037 – 3056

Easton R D, Srinivas K, Greene A J, 1997 "Do vision and haptics share common representations? Implicit and explicit memory within and between modalities", in *Journal of Experimental Psychology: Learning, Memory, and Cognition* **23** 153 – 163

Edelman S, Bülthoff H H, 1992 "Orientation dependence in the recognition of familiar and novel views of three-dimensional objects" *Vision Research* **32** 2385 – 2400

Edelman S, Weinshall D, 1991 "A self-organizing multiple-view representation of 3D objects" *Biological Cybernetics* **64** 209 – 219

Ernst M O, Lange C, Newell F N, 2007 "Multisensory recognition of actively explored objects", in *New Directions in Touch* Eds S Lederman, R L Klatzky *Canadian Journal of Experimental Psychology* **61** 242 – 253

Gauthier I, Hayward W G, Tarr M J, Anderson A W, Skudlarski P, Gore J C, 2002 "BOLD activity during mental rotation and viewpoint-dependent object recognition" *Neuron* **34** 161 – 171

Green D M, Swets J A, 1966 *Signal Detection Theory and Psychophysics* (New York: Wiley)

Hayward W G, Williams P, 2000 "Viewpoint dependence and object discriminability" *Psychological Science* **11** 7 – 12

Heller M A, Bracket D D, Salik S S, Scroggs E, Green S, 2003 "Objects, raised lines, and the haptic horizontal – vertical illusion" *Quarterly Journal of Experimental Psychology* **56** 891 – 907
James T W, Humphreys G K, Gati J S, Servos P, Menon R S, Goodale M A, 2002 "Haptic study of three-dimensional objects activates extrastriate visual areas" *Neuropsychologia* **40** 1706 – 1714
Kappers A M, Viergever R F, 2006 "Hand orientation is insufficiently compensated for in haptic spatial perception" *Experimental Brain Research* **173** 407 – 414
Kilgour A R, Lederman S J, 2006 "A haptic face-inversion effect" *Perception* **35** 921 – 931
Klatzky R L, Lederman S J, 1995 "Identifying objects from a haptic glance" *Perception & Psychophysics* **57** 1111 – 1123
Lederman S J, Klatzky R L, 1987 "Hand movements: a window into haptic object recognition" *Cognitive Psychology* **19** 342 – 368
Logothetis N K, Pauls J, 1995 "Psychophysical and physiological evidence for viewer-centered object representations in the primate" *Cerebral Cortex* **5** 270 – 288
Malach R, Reppas J B, Benson R R, Kwong K K, Jiang H, Kennedy W A, Ledden P J, Brady T J, Rosen B R, Tootell R B, 1995 "Object-related activity revealed by functional magnetic resonance imaging in human occipital cortex" *Proceedings of the National Academy of Sciences of the USA* **92** 8135 – 8139
Newell F N, Ernst M O, Tjan B S, Bülthoff H H, 2001 "Viewpoint dependence in visual and haptic object recognition" *Psychological Science* **12** 37 – 42
Newell F N, Findlay J M, 1997 "The effect of depth rotation on object identification" *Perception* **26** 1231 – 1257
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** 342 – 351
Palmer S, Rosch E, Chase P, 1981 "Canonical perspective and the perception of objects", in *Attention and Performance IX* Eds J B Long, A D Baddeley (Hillsdale, NJ: Lawrence Erlbaum Associates) pp 135 – 151
Peissig J J, Tarr M J, 2007 "Visual object recognition: do we know more now than we did 20 years ago?" *Annual Review of Psychology* **58** 75 – 96
Perrett D I, Harries M H, Looker S, 1992 "Use of preferential inspection to define the viewing sphere and characteristic views of an arbitrary machined tool part" *Perception* **21** 497 – 515
Reales J M, Ballesteros J, 1999 "Implicit and explicit memory for visual and haptic objects: Cross-modal priming depends on structural descriptions" *Journal of Experimental Psychology: Learning, Memory, and Cognition* **118** 219 – 235
Tarr M J, Bülthoff H H, 1998 "Image-based object recognition in man, monkey, and machine" *Cognition* **67** 1 – 20
Ullman S, 1998 "Three-dimensional object recognition based on the combination of views" *Cognition* **67** 21 – 44
Valentine T, 1988 "Upside-down faces: a review of the effect of inversion upon face recognition" *British Journal of Psychology* **79** 471 – 491
Verfaillie K, Boutsen L, 1995 "A corpus of 714 full-color images of depth-related objects" *Perception & Psychophysics* **57** 925 – 961
Wallis G, Bülthoff H H, 2001 "Effects of temporal association on recognition memory" *Proceedings of the National Academy of Sciences of the USA* **98** 4800 – 4804

ISSN 0301-0066 (print) ISSN 1468-4233 (electronic)

PERCEPTION

VOLUME 37 2008

www.perceptionweb.com