Haptic Own-Face Recognition

Sarah J. Casey and Fiona Newell

Department of Psychology, Trinity College Dublin, Ireland.
{caseys, fiona.newell} @tcd.ie

Abstract. This study investigated participants' ability to recognise a model of their own face from a series of distractor faces by touch. It also sought to discover whether or not an optimal stimulus orientation existed for the identification of target faces. Vision was found to differ significantly from each of the three haptic conditions (facing away from participant, facing towards participant, and learning, $p < 0.01$). Analysis of the three haptic conditions returned a main effect for recognition performance ($p < 0.05$), with performance differing significantly between away and learning conditions ($p < 0.05$), indicating the possible use of a feature-comparison search strategy. A greater number of target faces were correctly identified when faces were oriented towards the participants, contrary to the orientation in which a haptic representation of one's own face is naturally generated. This suggests that individuals may use a stored visual representation of their face when engaging in own-face haptic recognition.

1 Introduction

Over the last number of years there has been a surge in interest as to how information from different modalities is integrated to form coherent perceptions of objects in the world. Despite these advances, face recognition has been studied almost exclusively within the visual domain. The question thus arises as to whether the haptic system processes faces in the same manner as the visual system, or if face processing is modality specific. Kilgour and Lederman [1] have sought to readdress the balance in this respect, and examined the ability to identify and represent faces by hand. Participants were required to learn unfamiliar live human faces either visually, haptically or using both senses, and then try to recognise these target faces from among a group of distractor faces either within or across modalities. It was found that haptic input was more than sufficient for face identification. When the faces were explored both visually and by hand, recognition accuracy did not improve relative to touch alone. A further condition was tested in which participants learned and were then tested on recognition of masks of human faces. In this instance matching accuracy was found to decline compared to that with live faces, indicating that material properties of the faces such as texture play an important subsidiary role to that played by global properties in unfamiliar face haptic recognition.

As individuals have a vast amount of experience in recognising their own faces visually, either in the mirror or in photographic images, this experiment set out to investigate whether individuals could recognise their own face by touch. Does a common multi-modal representation exist wherein sensory input is shared between the haptic and visual systems?

We also sought to identify whether or not an optimal orientation exists for haptic recognition of one's own face. Previously, the examination of viewpoint dependency in visual and haptic object recognition found that, as in vision, haptic recognition of objects is much better when the study and test views are the same, that is, where as the visual system prefers the front view of objects best, the haptic system optimally recognises objects from the back [2]. It was also discovered that the transfer of object information between the visual and haptic modalities is view-point specific, with recognition performance being best when the objects were rotated 180° between learning in the first modality and recognition in the second. Here it was hypothesised that, like for 3-dimensional non-face objects, recognition would be better when a model of an individual's face was presented facing away from him or her, as this is the only view of his or her face that will have been experienced through touch.

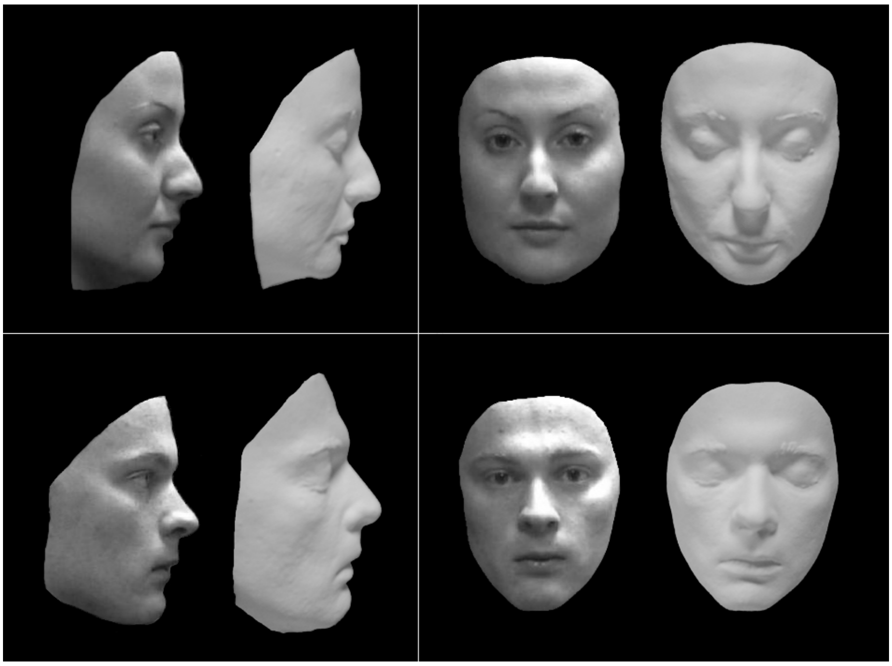

Figure 1. A sample of the face stimuli used in this study. On the top row from left to right is a female profile and the corresponding profile view of the stimulus face, and a female portrait with the corresponding portrait view of the stimulus face. On the bottom row is the same detail for a male face

2 Method

2.1 Participants

A mixture of 32 undergraduate and graduate students (16 male, 16 female, mean age 22.88 years, SD = 4.78) participated in this study. Testing procedures met with the standards delineated by the Ethics Committee of the Department of Psychology, Trinity College Dublin. Undergraduates received research credit as part of course requirements.

2.2 Materials

We made a plaster model of each participant's face as stimuli. Hair was covered with a swimming-cap and a thin layer of petroleum jelly was applied to the eyebrows and eyelashes. A thick layer of Accu-cast 880^{TM}, an alginate casting substance, was applied to the face and left to set for 8 minutes. As alginate is flexible even when set, a layer of gypsona plaster of paris bandages was applied over it and allowed to harden for 10 minutes. The mould was then removed from the participant's face. A thin layer of casting plaster was poured into each mould and allowed to stand for 20 minutes before a piece of plaster-soaked burlap was inserted and the remaining plaster poured in. The mount was inserted into the base of the mould - a square wooden pole (10x1x1 inches) which stood on a base measuring 6x6 x1inches. Each face was left to dry in the mould for 24 hours and when removed was allowed to dry for a further 24-48 hours before small imperfections were corrected using very fine-grain sandpaper.

2.3 Design

A repeated measures design was implemented and data were analysed with a Friedman analysis of variance (ANOVA, $\alpha = 0.05$).

2.4 Procedure

Approximately one week after the casting process, participants were recalled to complete a face recognition task. Each participant was instructed to select his or her own face from a line-up of eight faces, recognition to be carried out by touch. Participants could feel each stimulus for an unrestricted period of time, but were not permitted to return to any previously explored faces once they had moved on. There were four blocks in the experiment; participants were blindfolded for the first two. The orientation of the face stimuli in the first two blocks was counterbalanced across participants. The presentation order of the target face among distractor stimuli was randomised across blocks for each participant. The blindfold was removed for block

3, in which the face stimuli always faced away from the participant to obscure shape from view. In this block, before feeling each stimulus, participants were asked to feel their own face for a full 10 seconds. In the final block, participants were required to visually identify their own face models. Participants identified their faces by giving the location of their face model in the 'line-up'.

3 Results

The mean percentage correct recognition rates are presented in Figure 2. Data were analysed using a Friedman analysis of variance (ANOVA), with the alpha level set at 0.05. A main effect was found for condition indicating differential performance among the four levels ($\chi^2 = 21.35$, df = 3, $p < 0.0001$). Post-hoc analysis indicated that the performance on the visual recognition task was significantly different from performance in the other three conditions ($p < 0.01$). Results of the Wilcoxon Matched-Pairs tests can be found in Table 1. A second Friedman ANOVA was run, excluding the vision condition. A main effect was returned for these haptic conditions ($\chi^2 = 8.45$, df = 2, $p < 0.02$). The Wilcoxon Matched-Pairs test shows that the second ANOVA result is due to a significant difference between the facing away and the learning conditions ($p < 0.03$). As the learning condition was always the last

Table 1. Results of the Wilcoxon Matched-Pairs tests, N = 32

Condition	T	Z	p-level
Towards * Away	11.0	1.68	0.09
Towards * Learning	16.5	1.12	0.26
Away * Learning	25.5	2.20	0.03*
Vision * Towards	8.5	3.07	0.01**
Vision * Away	11.5	3.73	0.001***
Vision * Learning	0.0001	2.80	0.01**

* Significant at the 0.05 level, ** Significant at the 0.01 level, *** Significant at the 0.001 level.

haptic condition presented, a final ANOVA was conducted to rule out an order effect. When analysed by block, no significant differences emerged (($\chi^2 = 5.45$, df = 2, $p < 0.07$). Performance in each condition was compared to that which might be expected by chance. Participants were found to perform significantly better than chance in all conditions ($p < 0.05$), except for the facing away condition ($p = 0.13$).

4 Discussion

It is clear that for own-face recognition visual performance is far superior to that in any of the haptic conditions. Firstly, individuals may simply have greater visual expertise, thus encoding more robust visual representations of their faces. Secondly, live faces comprise both geometric and material properties, both of which are necessary in the derivation of haptic representations; accordingly it was found that, matching accuracy declined when participants were required to haptically recognise

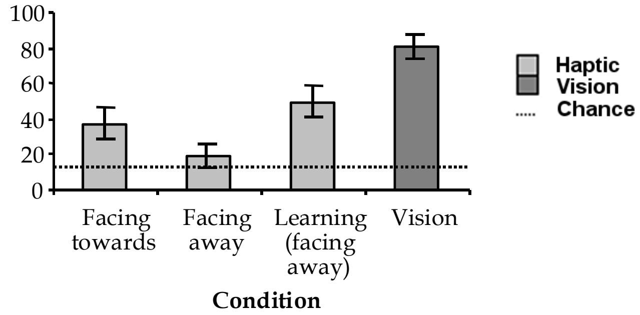

Fig 2. Mean percentage of correct recognition responses

masks of faces, presumably due to the absence of material properties [1]. Similarly, the face models employed in this study provide very little material information, thus hindering greater accuracy in haptic face recognition. It is interesting to note that although visual self-recognition was more accurate, it did not reach ceiling level, possibly due to missing characteristics such as pigmentation [3].

Participants performed significantly above chance in all conditions except when models were presented facing away from the participants. In both the facing away and facing conditions, it is presumed that face recognition is facilitated by stored haptic representations as no learning time is allowed. Insufficient information is extracted from the face model to allow above chance performance in the facing away condition, suggesting that something else is required to carry the representation over threshold for a positive identification. There are two possible mediators: a feature comparison strategy, and visual imagery. Performance in the learning condition was significantly better than in the facing away condition; models in both conditions were oriented in the same direction. The increased positive recognition scores in the learning condition may have been facilitated by the refreshing of haptic representations of participant's own faces before each trial. This condition could also have allowed participants to develop a comparison strategy, whereby they could actively compare their own features to those of the distractor faces. Haptic face representations may not be newly generated in the learning condition, because when face models were oriented towards the participants, a condition in which no learning time was allowed, performance was also greater than chance. The higher percentage of positive recognition responses in the learning condition could not be attributed to practice, as there was no evidence of an order effect in the data.

Although no effect of stimulus orientation was obtained, the trend apparent in the data is somewhat surprising. On average there were more positive identifications in the facing towards participant condition. This is unusual in terms of previous findings for 3-dimensional objects where haptic recognition of objects was found to be viewpoint-dependent [2]. It was expected that participants would find recognition easier when faces were oriented away from them, as this corresponds to the natural haptic learning 'view' for one's own face. Could it be that participants are using some form of mental imagery in this task, which elevates performance above chance levels? Although Kilgour and Lederman [1] did not find any evidence supporting the use of visual imagery in their study with unfamiliar live faces their participants generated representations of the stimuli in a controlled fashion. However, participants in this study already had prior visual knowledge of their own faces, and may have relied on a more robust visual representation gained from photographs or video footage, and daily use of mirrors. Although face representations generated from mirror usage are reversed, it is the most commonly associated visual experience when touching one's own face. It is possible that given the frequency of these associated experiences, individuals have become quite adept at translating these visual representations into haptic ones.

5 Acknowledgements

The authors wish to thank Tanja Khosrawi for her dedicated assistance in creating the stimuli faces.

References

1. Kilgour, A. R. and Lederman, S. J. Face recognition by hand. Perception & Psychophysics, 64 (3), 339-352 (2002).
2. Newell, F. N., Ernst, M.O., Tjan, B. S., and Bülthoff, H. H. Viewpoint dependence in visual and haptic object recognition. Psychological Science, 12 (1), 37-42 (2001).
3. Bruce, V., Healey, P., Burton, A. M., Doyle, T., Coombes, A., and Linney, A. Recognising facial surfaces. Perception, 20, 755-69 (1991)