Cognition 165 (2017) 97–104

Contents lists available at ScienceDirect

Cognition

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

Original Articles

Camera-to-subject distance affects face configuration and perceived identity

Eilidh Noyes, Rob Jenkins *

Department of Psychology, University of York, UK

ARTICLE INFO

Article history:
Received 13 December 2016
Revised 1 May 2017
Accepted 10 May 2017
Available online 17 May 2017

Keywords:
Face recognition
Face matching
Identification
Configural processing
Camera
Passport

ABSTRACT

Face identification is reliable for viewers who are familiar with the face, and unreliable for viewers who are not. One account of this contrast is that people become good at recognising a face by learning its configuration—the specific pattern of feature-to-feature measurements. In practice, these measurements differ across photos of the same face because objects appear more flat or convex depending on their distance from the camera. Here we connect this optical understanding to face configuration and identification accuracy. Changing camera-to-subject distance (0.32 m versus 2.70 m) impaired perceptual matching of unfamiliar faces, even though the images were presented at the same size. Familiar face matching was accurate across conditions. Reinstating valid distance cues mitigated the performance cost, suggesting that perceptual constancy compensates for distance-related changes in optical face shape. Acknowledging these distance effects could reduce identification errors in applied settings such as passport control.

1. Introduction

Recognising a person's face involves mapping an image onto an identity (Bruce & Young, 1986). Completing this mapping reliably is a challenge, because images of the same face can be as varied as images of different faces (Jenkins, White, Van Montfort, & Burton, 2011). One way of characterising this challenge is as a signal-to-noise problem: the signal is the 'true' appearance of the person's face, the noise consists of deviations from that appearance, and the task of the visual system is to extract the signal from the noise.

For some time, the search for an identity signal has centred on *configural* information—the idiosyncratic spatial layout of a person's face, typically defined in terms of 'metric distances between features' (Tanaka & Gordon, 2011). The proposal is that each face has a unique configuration, which viewers come to learn, and that knowledge of that configuration allows the viewer to recognise that particular person (Diamond & Carey, 1986; Maurer, Le Grand, & Mondloch, 2002). A parallel research effort has documented effects of different types of noise on identification accuracy. For example, many studies have measured the impact of viewing angle, facial expression, and lighting conditions on observers' performance (Bruce, 1982; Hill & Bruce, 1996; Johnston, Hill, & Carman, 1992; Troje & Bülthoff, 1996; Young, McWeeny, Hay, & Ellis, 1986).

Relatively few studies have examined effects of viewing *distance*. Those studies have typically been motivated by forensic questions concerning eyewitness testimony (De Jong, Wagenaar, Wolters, & Verstijnen, 2005; Greene & Fraser, 2002; Hahn, O'Toole, & Phillips, 2016; Lampinen, Erickson, Moore, & Hittson, 2014; Lindsay, Semmler, Weber, Brewer, & Lindsay, 2008; Loftus & Harley, 2005; Wagenaar & van der Schrier, 1996). As such, they have focused almost exclusively on two inter-related questions, (i) What is the maximum distance at which we can recognise a face?, and (ii) What is the minimum information required to recognise a face? These questions arise directly from the optics of the situation: more distant objects project smaller retinal images (with implications for spatial frequency content; Loftus & Harley, 2005). But optics gives us another reason to take viewing distance seriously: changes in viewing distance affect configural information in the face image.

Perhaps the most compelling demonstration of this configural change comes not from face recognition research, but from analyses of perspective in portraiture. Harper and Latto (2001) photographed models' faces at different camera-to-subject distances (0.32 m, 0.71 m, 2.70 m), and rescaled the faces to the same interocular distance. As Fig. 1 illustrates, faces look convex when close, and flatter from afar. In other words, the same face appears to have

* Corresponding author at: Department of Psychology, University of York, York YO10 5DD, UK.

E-mail address: rob.jenkins@york.ac.uk (R. Jenkins).

http://dx.doi.org/10.1016/j.cognition.2017.05.012
0010-0277/

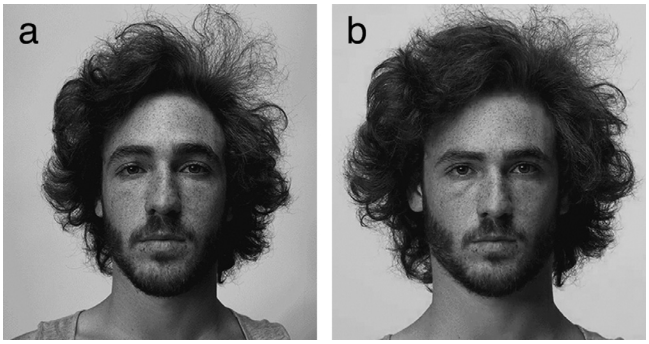

Fig. 1. Two photographs of the same face taken from different viewing distances: (a) ∼0.20 m; (b) ∼3.00 m. Photos are shown rescaled to the same interocular distance. © Dan Vojtěch 2016. Reproduced with permission.

quite distinct shapes when viewed from different distances. Indeed, participants in Harper and Latto's (2001) study gave higher weight estimates for the models as camera-to-subject distance increased.

More recent work (Bryan, Perona, & Adolphs, 2012) has shown that social inferences from faces also change with camera-to-subject distance. Viewers' ratings of trustworthiness, competence, and attractiveness were all lower for photos that were taken closer (0.45 m) than for photos that were taken further away (1.35 m), presumably because such inferences rely partly on shape cues (Oosterhof & Todorov, 2008; Zebrowitz, 2011).

Taken together, these studies confirm that (i) face shape in the image changes with camera-to-subject distance, and (ii) these shape changes are large enough to have psychological consequences, even when the difference in viewing distance is small (e.g. 1–2 m). Despite the clarity of these findings, there has been almost no attempt to pursue their implications for the important issue of face identification (see Liu, 2003, for an exception). This oversight is perhaps surprising, given the emphasis on configural information in the face recognition literature. If viewing distance alters configural information, and configural information is key to face identification, it follows that viewing distance should affect face identification. That is the argument that we examine in the present studies. We begin with a direct test of the first premise, that viewing distance alters configural information.

2. Study 1. Camera-to-subject distance affects feature-to-feature measurements

The purpose of this study was to relate changes in camera-to-subject distance to changes in facial configuration. The apparent size of an object clearly varies with viewing distance, in the sense that the size of the retinal image changes. Linear changes in the size of a 2D face image (as when a photograph is rescaled) do not affect configural layout because they do not affect the relative distances between features. Consistent with the conservation of configural layout over size changes, behavioural and neuroimaging studies have found that face recognition is unaffected by linear rescaling (Andrews & Ewbank, 2004; Bindemann, Burton, Leuthold, & Schweinberger, 2008; Bruce, Burton, Carson, Hanna, & Mason, 1994; Grill-Spector et al. 1999). For 3D objects (e.g. live faces as opposed to face photographs), the optical situation is very different. Changes in camera-to-subject distance produce non-linear changes in the image, such that different parts of the image are affected to differing degrees (Latto & Harper, 2007; Pirenne, 1970). For convex objects such as faces, distant viewing leads to flatter appearance, whereas closer viewing leads to more convex appearance (see Fig. 1). To tie this optical transformation directly to the notion of configuration in the face perception literature, we measured distances between key facial features in photos that were taken at different viewing distances. The expectation was that, as a reflection of the flat-to-convex transformation, measures nearer the edge of the face would be compressed relative to measures nearer the centre of the face.

2.1. Photographic procedure

The images used for all of these studies were face photographs of 18 consenting undergraduates at the University of York. These volunteer models were photographed in two separate sessions, one week apart. In each session, each model was photographed at two distances—*Near* (camera-to-subject distance = 0.32 m) and *Far* (camera-to-subject distance = 2.70 m), following Harper and Latto (2001). This regime resulted in four photographs for each of the 18 models: Week 1 *Near*, Week 1 *Far*, Week 2 *Near*, and Week 2 *Far* (72 photos in total). All models were photographed with a neutral expression using an Apple iPhone 5s on default settings. Photos were then cropped around the head to remove extraneous background. For anthropometric analysis, each image was scaled to an interocular distance of 150 pixels, with aspect ratio preserved.

2.2. Anthropometric analysis

We follow Burton, Schweinberger, Jenkins, and Kaufmann (2015) in extracting from the literature those feature-to-feature distances that have been offered as specific examples: distance between the corner of the eye and the edge of the nose (left and right; Leder & Carbon, 2006), distance between the corner of the nose and the corner of the mouth (left and right; Leder & Bruce, 2000), and distance between the nose and the mouth (Leder & Carbon, 2006; see Burton et al., 2015, for precise anatomical definitions). This resulted in five measurements in total for each photograph, which were made using the Ruler tool in Adobe PhotoShop. Fig. 2 shows these five measurements for *Near* and *Far* photos of one volunteer model.

2.3. Results and discussion

For each of the five feature-to-feature metrics, we conducted a 2 × 2 repeated-measures ANOVA with the factors of *Photographic Session* (*Week 1* versus *Week 2*) and *Camera-to-Subject Distance* (*Near* versus *Far*). Results of these analyses are summarised in Table 1.

As can be seen from Table 1, *Photographic Session* had no significant effect on any of the measurements ($p > 0.1$ for all), indicating that incidental changes in viewpoint and expression for *Week 1* versus *Week 2* were negligible. In this context, *Camera-to-Subject Distance* systematically affected some measures but not others. The relatively peripheral nose-to-mouth measurements were larger for *Far* images than for *Near* images, whereas the more central eye-to-nose measurements were statistically equivalent at the two camera distances we compared. This pattern in the anthropometric data corroborates the flatter appearance of the *Far* images and the more convex appearance of the *Near* images, and is consistent with the differential weight estimates in previous studies (Harper & Latto, 2001). More importantly for the current study, it confirms the non-linear effect of camera-to-subject distance on configural information: some feature-to-feature measurements changed substantially and others did not (see Smith, 2016, for a computational perspective). We next used a paired matching task to assess the implications of these configural changes for perception of facial identity.