Perception, 2014, volume 43, pages 90–98

doi:10.1068/p7436

SHORT REPORT

Looking at a blurry old family photo? Zoom out!

Kimeya Shahangian, Ipek Oruc§
Department of Ophthalmology and Visual Sciences, University of British Columbia,
818 West 10th Avenue, Vancouver, BC V5Z 1M9, Canada; e-mail: ipor@mail.ubc.ca
Received 18 December 2012, in revised form 17 January 2014

Abstract. We investigated recognition of blurry faces and whether viewing size affects identification of such severely degraded images. Despite the common belief that face perception relies on middle spatial frequencies, the critical spatial frequency band for face recognition is not fixed but rather depends on size. This is especially pronounced at small sizes, where observers choose to utilize lower, rather than middle, frequencies to identify a face. Here we assessed recognition of identity via a novel use of the face adaptation paradigm. We examined face identity aftereffects of blurry and intact adaptors at two sizes. Intact adaptors induced significant aftereffects regardless of size. Small, but not large, blurry adaptors produced aftereffects despite the fact that both contained exactly the same level of facial detail. This suggests an inability to utilize low-frequency information for perceiving identity in large faces. We conclude that (1) size is a key factor in human face recognition processes and (2) coarse facial images are better recognized at small sizes.

Keywords: face recognition, adaptation, aftereffects, scale invariance, effects of size, optimal viewing distance

1 Introduction

The spatial frequencies used by human observers to recognize[(1)] faces have been studied extensively in the past. Despite the current consensus in the field that face recognition is based on middle spatial frequencies,[(2)] the reports of numerous studies on the topic have been highly diverse. In addition to studies consistent with the utilization of middle frequencies (Costen, Parker, & Craw, 1996; Fiorentini, Maffei, & Sandini, 1983; Gaspar, Sekuler, & Bennett, 2008; Goffaux, van Zon, & Schiltz, 2011; Gold, Bennett, & Sekuler, 1999; Näsänen, 1999; Parker & Costen, 1999; Schyns, Bonnar, & Gosselin, 2002), some reported high frequencies (Hayes, Morrone, & Burr, 1986) and others found very low spatial frequencies of around 1–3 cycles per face-width to be highly effective for face recognition (Harmon, 1973; Rubin & Siegel, 1984; Sinha, 2002a, 2002b). It has been suggested that this apparent discrepancy between studies may stem from differences in experimental methodologies (Gold et al., 1999). Although different spatial frequencies have been empirically observed to be critical for identification of different-size face stimuli in several studies (Näsänen, 1999; Willenbockel et al., 2010) as well as for recognition of facial expressions across sizes (Smith & Schyns, 2009), generally speaking, face size has not been considered to play a critical role in face recognition.

A recent study systematically varied face size in a critical-band masking paradigm and found that spatial frequencies used by human observers are strongly influenced by viewing size (Oruc & Barton, 2010a). This effect is especially pronounced at smaller sizes. While spatial frequencies in the middle range (around 8 cycles per face-width) were dominant for

§ Corresponding author.

(1) From here on, the term *recognition* refers to the ability to individuate an exemplar rather than to indicate if it was seen before—an alternative convention.

(2) Throughout, *spatial frequency* refers to frequency in object units—specifically, cycles per face-width (as opposed to absolute frequencies in cycles per degree) unless otherwise specified.

recognizing faces that were about 5 deg or larger, critical frequencies shifted to lower bands for smaller faces, gradually approaching 3 cycles per face-width as the face size was reduced below 2 deg (see figure 2). This study therefore suggested that the commonly accepted rule of thumb that 'faces are recognized using middle frequencies' must be qualified with the size of the face. In particular, small faces are recognized using low spatial frequencies—that is, gross forms in the image. It is, however, unclear whether this size dependency truly characterizes face recognition in the real-life context and has any practical relevance. First off, these results are based on low-contrast near-threshold face images viewed in visual noise. Suprathreshold perception does not always follow subthreshold characteristics. In addition, the methodology used in that study—critical-band masking—yields peak spatial frequencies; that is, those that are maximally diagnostic of identity. This does not, however, imply that these spatial frequencies are used exclusively, or that other frequency bands are inaccessible for the purpose of identification.

In the present study we examine human observers' ability and preference for utilizing low spatial frequency bands in small versus large faces viewed in easily visible suprathreshold contrast conditions. Specifically, we hypothesize that faces are recognized through low spatial frequencies in small sizes. However, for large faces, observers shift to higher frequency bands and, importantly, lose their ability to access lower frequency bands (relative to the whole object—eg, in cycles per face-width) for the purpose of recognition. Consider the blurred face images in figure 1. Looking at the large image at the top, most observers report seeing a female face but are unsure of its identity. Surprisingly, viewing the same image at a small size (bottom) immediately makes it clear that this is in fact Angelina Jolie. The same effect can be obtained by looking at the large face image from a distance of about 3–4 m. The fact that the small blurry face is readily identifiable shows that the level of facial detail remaining in the blurry image is sufficient for facial recognition. These facial details do not become invisible at the large size, yet the human visual system can no longer utilize them for the purpose of recognition. This illustration demonstrates that face recognition processes are dependent on size at a fundamental level. Here we use a face adaptation paradigm to psychophysically demonstrate and quantify this observation.

Perception of faces is influenced by prior exposure to other faces in a manner that biases away from the attributes of the previously seen face. For example, viewing a male face causes a subsequently seen face with ambiguous gender to be perceived as female (eg Webster, Kaping, Mizokami, & Duhamel, 2004). This bias, called the face adaptation aftereffect, has been observed for a number of facial attributes such as identity, expression, and gaze direction (Benton et al., 2007; Benton, Jennings, & Chatting, 2006; Fox, Oruc, & Barton, 2008; Jenkins, Beaver, & Calder, 2006; Jiang, Blanz, & O'Toole, 2006; Leopold, O'Toole, Vetter, & Blanz, 2001; Webster et al., 2004). In the present study we measure identity aftereffects on faces produced by a morphing technique that blends together two distinct female faces. Under normal conditions, images from the middle portion of the morph series are perceived to be of ambiguous identity, looking equally like either of the constituent faces. This is expected, as the morph faces in the middle portion are made of roughly equal parts of the two identities. However, if observers are shown one of the constituent faces for a few seconds prior to viewing the morph test face, then they distinctly perceive the identity of the other constituent face in this ambiguous stimulus.

In the present study we utilize this adaptation procedure to indirectly deduce whether the observer can recognize the identity of a blurry face. We reason that identity after-effects can occur only if observers can recognize the adapting face. Thus, we gauge recognition of the adapting face from the magnitude of the identity aftereffect it can produce. We measure aftereffects of viewing large and small adapting faces on medium-sized test faces.

Figure 1. Most viewers are unable to identify the large blurry face seen at the top. The same image is more recognizable when viewed in a smaller size. The reader can confirm that this is not an image artifact by viewing the large image from a distance of 3–4 m.

Our adapting faces are either intact or blurred, making up the following four adapting conditions: (1) large intact, (2) large blurred, (3) small intact, and (4) small blurred. We based the choice of cut-off frequency for our blurry images on the Oruc and Barton (2010a) study. In figure 2 we show peak spatial frequency for face recognition as a function of face size (figure replotted with permission from Oruc and Barton, 2010a). In the blurry adapting images we aimed to remove the spatial frequencies that are utilized at the large condition (10 deg) but retain the band utilized in the small condition (2 deg). For this purpose, we low-pass filtered the images at a cut-off frequency of 6 cycles per face-width (dashed arrow), corresponding to 0.6 cpd and 3 cpd at the large and small sizes, respectively. Although the blurry image filtered in this manner contains the peak frequencies utilized at small viewing conditions but not those utilized at large viewing conditions according to Oruc and Barton (2010a), it nevertheless contains an identical level of facial detail regardless of the size. All test stimuli are medium sized and intact. If the identity of the blurry adapting face is discernible only at the small size, but not at the large size, we predict that we will find significant aftereffects due to small-blurred adaptors, but no or reduced aftereffects with large-blurred adaptors. The intact adaptors serve as a controls.

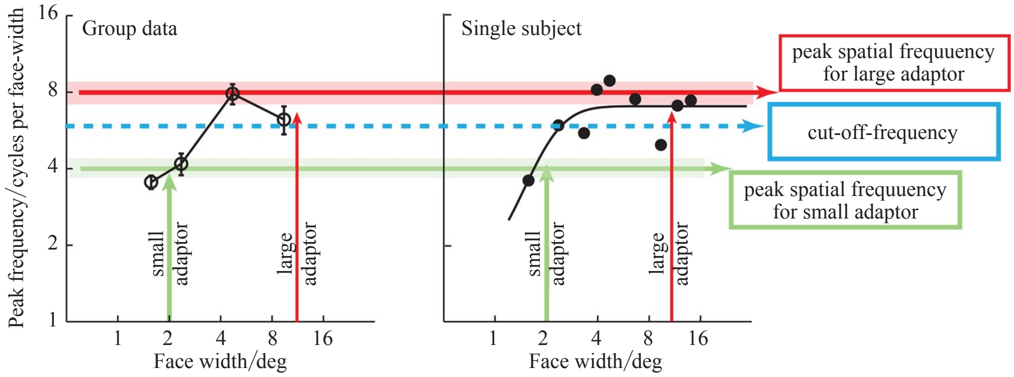

Figure 2. [In color online at http://dx.doi.org/10.1068/p7436] Illustration of the rationale for the choice of cut-off frequency in the generation of the blurry adapting stimuli. Peak spatial frequencies utilized for face recognition are shown as a function of face size (data reproduced from Oruc and Barton, 2010a). This shows that recognition of small faces (<5 deg in width) depends on a lower spatial frequency band within the object spectrum than that utilized for larger faces. Our small adaptor was 2 deg in width corresponding to a band of around 4 cycles per face-width (green arrow). Our large adaptor was 10 deg in width, utilizing a band of around 8 cycles per face-width (red arrow). In preparation of our blurry stimuli, our aim was to remove the frequency band (red) optimal for the large size yet retain the band (green) optimal for the small size. Thus we used a cut-off frequency of 6 cycles per face-width (dashed, cyan).

2 Results

Figure 3 shows face identity aftereffects for all four adapting conditions: large intact, small intact, large blurry, and small blurry. Aftereffect magnitude was unaffected by adaptor size ($F_{1,14} = 2.41$, $p = 0.14$), but there was a main effect of image condition ($F_{1,14} = 24.73$, $p < 0.001$) that was explained by a significant interaction between size and image condition ($F_{1,14} = 5.22$, $p = 0.038$). Aftereffect magnitude was significantly larger than zero for all conditions (large intact $M = 16\%$, small intact $M = 12\%$, small blurry $M = 8\%$; t-test, all ps<0.05) except for the large blurry condition, which failed to show any aftereffect

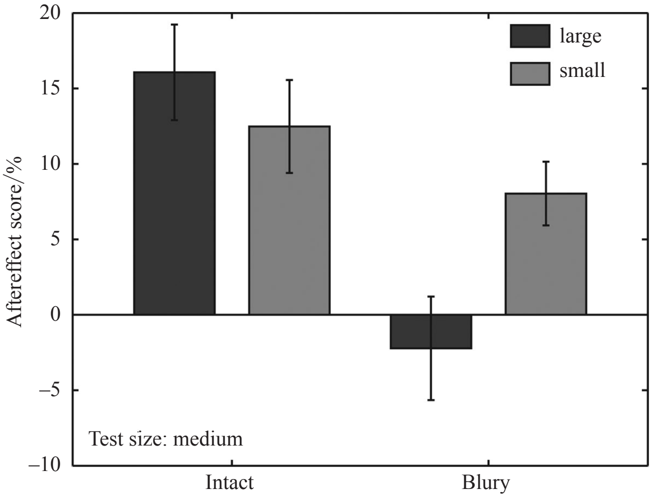

Figure 3. Experimental results. Aftereffect magnitude is shown for the four adapting conditions: large intact, small intact, large blurry, and small blurry. Both intact adaptors produced significant aftereffects. Aftereffect size was slightly reduced for the small-intact condition compared with the large-intact condition, though this difference was not statistically significant. In the blurry condition this pattern reversed: aftereffect size was significantly larger in the small-blurry condition compared with the large-blurry condition, which failed to generate any aftereffect.

($M = -2\%$, $p > 0.5$). A posteriori pairwise comparisons showed a significant difference between the large blurry condition and both large and small intact conditions (Tukey–Kramer, all *p*s < 0.05), while none of the remaining conditions differed significantly from each other. (See figure S3 in the supplementary material at http://dx.doi.org/10.1068/p7436 for response curves as a function of morph level and adapting face.)

3 Discussion

Overall, intact adapting faces at both the small and the large sizes generated significant and comparable aftereffects on the perception of medium-sized intact test faces. The magnitudes of these effects are consistent with previous studies that measured face identity aftereffects under similar experimental settings (eg Oruc and Barton, 2010b, experiment 2). In addition, substantial aftereffects were obtained with small blurry adapting faces, which were numerically, but not statistically, lower than the intact adaptors. The large blurry adapting faces, on the other hand, failed to generate aftereffects. The aftereffect magnitude in the small intact condition was slightly lower than the large intact condition, even though the size mismatch between these and the test were approximately equal. This decrease may be due to our scaling manipulation in which the small intact image was formed by subsampling the large image, and thus contained fewer pixels. This minor methodological shortcoming did not impact our main conclusion—rather, it emphasized our critical finding in the blurry condition where the direction of change was in the opposite direction.

It is important to note that the large blurry adaptors contained the same physical information and level of detail regarding facial identity. So why did they fail to produce aftereffects? One possibility is a simple effect of size. It is known that aftereffects are reduced due to size mismatches between the test and adapting faces. So perhaps it is the large adapting size in general that is ineffective. However, the intact adaptor condition rules out this explanation: the large intact adaptor generates as large, if not larger, an effect as the small intact adaptor. The lack of an aftereffect in the large blurry condition is rather due to the interaction between size and utilization of spatial frequency content. Although sufficient identity information is available in the blurry image, as evidenced by the small-blurry results, this information becomes inaccessible at the large size.

Previous studies that examined recognition of various visual objects, such as letters, have found that optimal spatial frequencies in object units change with stimulus size (Chung, Legge, & Tjan, 2002; Majaj, Pelli, Kurshan, & Palomares, 2002; Oruc & Barton, 2010a). Chung et al. (2002) suggested that the observed peak spatial frequencies were based jointly on the diagnosticity of object frequency bands and the contrast sensitivity corresponding to that band at the given size, biasing peak recognition frequencies in the direction of peak contrast sensitivity. One study later examined this idea by introducing white noise at levels high enough to drown out the equivalent internal noise, thus flattening the effective contrast sensitivity curve (Oruc & Landy, 2009). Under these conditions, switching to other frequency bands no longer confers any advantage, yet human observers continued to behave as before. While this result was not supportive of the CSF-based account of the Chung et al. (2002) study, it also could not refute it completely since channel switching may be a long-term adaptation to the shape of the CSF, unlikely to change in response to the brief presence of the experimental white noise.

Can the CSF-based account explain the results presented here? The six cycles per face-width cut-off frequency corresponded to 3 cpd at the small adapting size (2 deg), compared with 0.6 cpd at the large adapting size (10 deg). Certainly, detection contrast thresholds would be slightly lower at 3 cpd; in other words, a 3 cpd pattern just visible at the small size may in fact be rendered invisible if viewed at the large size, at 0.6 cpd. However, this does not apply at the suprathreshold viewing conditions of the present study. Information regarding facial identity at 3 cpd that is easily visible in the small size at suprathreshold contrast is also fully visible at the large size at 0.6 cpd. Generally speaking, visibility is improved by enlarging a suprathreshold image, because at high-contrast levels it is limited predominantly by resolution constraints rather than contrast sensitivity. The level of detail in our blurry faces that enables recognition at the small size remains available and visible at the large size. Thus a CSF-based account is unlikely to explain why recognition of blurry faces is hampered at only the large size.

Our blurry images were low-pass filtered to remove all spatial frequencies beyond six cycles per face-width. With such severely degraded images, the real surprise is not that large sizes are ineffective adaptors but the fact that these blurry images can generate any aftereffects at all, as they do in the case of the small size. At this blur level fine details of the facial features are not discernible. It has been shown previously that human observers perform well with highly impoverished blurry faces (Burton, Wilson, Cowan, & Bruce, 1999; Sinha, 2002a, 2002b; Yip & Sinha, 2002). The present results suggest this ability may be limited to small sizes. The effect of size, or viewing distance, on efficiency of face recognition informs our current understanding of the mental processes behind face perception as well as real-life applications of face identification. It has been suggested before that faces are easier to recognize up close than far away (Loftus & Harley, 2005). Our results show that in conditions of limited visibility such as grainy, low-resolution images, which are more typical in the real life context, the effective strategy is the exact opposite: blurry faces are more recognizable at small sizes. A recommendation for real-life practice that immediately follows from this finding is that observers routinely dealing with blurry visual input, such as security personnel monitoring on-screen video feed of people at an airport, should view faces at small sizes for identification purposes.

4 Methods

4.1 *Subjects*

Fifteen subjects (seven females and eight males, ages 18–51 years) with normal or corrected-to-normal vision participated in this study. The protocol was approved by the review boards of the University of British Columbia and Vancouver Hospital, and informed consent was obtained in accordance with the principles in the Declaration of Helsinki.

4.2 *Experimental setup*

The experimental protocol was programmed and run using Superlab version 4.5 (http://www.superlab.com) on a Vaio Sony laptop (model PCG-71311L) with 15.6" screen. Subjects were seated approximately 62 cm from the computer screen in a dimly lit room. At this viewing distance the test images subtended 6.5 deg per face-width and the small and large adaptors were 2 and 10 deg per face-width, respectively.

4.3 *Stimuli*

Three female faces displaying a neutral expression were selected from the Karolinska Database of Emotional Faces (Lundqvist & Litton, 1998). Face images were converted to grayscale using Photoshop CS 8 (http://www.adobe.com). An oval aperture was superimposed on all images. Mean luminance and root-mean-squared contrast within the oval aperture was equalized across the three face images using in-house custom Matlab scripts (http://www.mathworks.com). Background luminance was set to half maximum luminance. For further details on stimulus generation see Oruc and Barton (2010b).

4.3.1 *Test stimuli.* We created morph series between all three pairs of faces using FantaMorph 4 (http://www.fantamorph.com), resulting in three distinct morph series. Each morph series contained 41 images that gradually blended the two constituent faces, F1 and F2, in 2.5% steps.

Only the middle 13 morph images ranging from 65% F1–35% F2 to 35% F1–65% F2 were used as test stimuli (supplementary figure S1c). The original constituent faces used to generate the morph stimuli served as adapting faces.

4.3.2 *Adapting stimuli.* While the test stimuli were always medium size (6.5 deg) and intact, adapting stimuli were one of two sizes, large (10 deg) or small (2 deg), and one of two image conditions, intact or blurry, making up the four adapting conditions. To generate the blurry images, intact adapting images were low-pass filtered at 6 cycles per face-width cut-off frequency using in-house custom Matlab scripts that implemented a Butterworth filter with the squared gain function, $G^2(f) = 1/[1 + (f/f_c)^{2n}]$, where f denotes spatial frequency and f_c denotes the cut-off frequency. The filter order, n, was set to 5 to ensure fairly steep attenuation while minimizing ringing artifacts.

Different size adaptors were obtained by first resizing the intact images in Photoshop using bicubic interpolation. The low-pass filter was applied after this step separately to each image size. We adopted this specific order of filtering after resizing to prevent any artifacts that may have been introduced by the resampling to contaminate the final stimulus image. To verify that the blurry adaptors had indeed identical spatial frequency content at the large and small sizes, we plot the amplitude spectrum as a function of spatial frequency averaged across orientations (figure S2). This comparison confirms that the blurry adaptors at the small and large sizes were nearly identical in their physical content.

4.4 *Experimental protocol*

A typical trial started with a 5s adapting period displaying one of two constituent faces, F1 or F2. This was followed by a noise (50 ms), fixation (150 ms), and blank (150 ms). Following that, one of the 13 morph images was presented to the participant for 300 ms. This was followed by a 150 ms blank and a choice screen displaying the two constituent faces. The participant's task was to indicate whether the test face resembled F1 or F2. The next trial started as soon as the participant entered their response by pressing one of two keys on the computer keypad. No feedback was provided. All 13 test images were presented twice, once each after viewing F1 and F2. For each morph series there were four adapting conditions—large intact, small intact, large blurry, and small blurry—resulting in 104 trials. There were three distinct morph series, resulting in a total of 312 trials. Each participant completed all 312 trials in a randomized order. The participants received a break after 52 trials to help them stay focused. They were able to choose when to continue after they had received the break.

Prior to starting the experimental trials, participants were familiarized with the three female faces that were used. First, participants were asked to freely view the face images and complete a 10 trial experiment where one of the faces was shown for 300 ms, and the participant had to determine which of the three faces it was. Upon answering at least 9 of these questions correctly, they completed a 48 trial training block in which the task was identical to the experimental protocol with the following differences: there was no adapting period, the test image was shown for a total of 1000 ms, and only the two outmost morphs (65% of each of the images) were shown as the test.

4.5 *Data analysis*

For convenience, a response of F1 was assigned a value of 0, and a response of F2 was assigned a value of 1, within each morph series. The sum of the score they achieved for the 13 morphs of the same face pair presented with the same adaptor was determined, yielding a maximum score of 13 (ie if they responded F2 on all 13 trials). We then determined the aftereffect for each condition for each face pair by finding the difference between the scores after adapting to F1 and after adapting to F2. Since adapting to F1 biases the observer to perceive F2 in the ambiguous test images, the sum of the responses is expected to be larger than that after adapting to F2; we define aftereffect score as the difference between the two. This value is expected to be a positive number when face identity aftereffects are observed. The aftereffect scores from each adapting condition were averaged for the three morph series for each participant.

Acknowledgment. This work was supported by NSERC Discovery Grant RGPIN 402654-11.

References

Benton, C. P., Etchells, P. J., Porter, G., Clark, A. P., Penton-Voak, I. S., & Nikolov, S. G. (2007). Turning the other cheek: the viewpoint dependence of facial expression after-effects. *Proceedings of the Royal Society of London B: Biological Sciences*, **274**(1622), 2131–2137.

Benton, C. P., Jennings, S. J., & Chatting, D. J. (2006). Viewpoint dependence in adaptation to facial identity. *Vision Research*, **46**, 3313–3325.

Burton, A. M., Wilson, S., Cowan, M., & Bruce, V. (1999). Face recognition in poor-quality video: Evidence from security surveillance. *Psychological Science*, **10**, 243–248.

Chung, S. T. L., Legge, G. E., & Tjan, B. S. (2002). Spatial-frequency characteristics of letter identification in central and peripheral vision. *Vision Research*, **42**, 2137–2152.

Costen, N., Parker, D., & Craw, I. (1996). Effects of high-pass and low-pass spatial filtering on face identification. *Perception & Psychophysics*, **58**, 602–612.

Fiorentini, A., Maffei, L., & Sandini, G. (1983). The role of high spatial frequencies in face perception. *Perception*, **12**, 195–201.

Fox, C. J., Oruc, I., & Barton, J. J. (2008). It doesn't matter how you feel. The facial identity aftereffect is invariant to changes in facial expression. *Journal of Vision*, **8**(3):11, 11–13.

Gaspar, C., Sekuler, A. B., & Bennett, P. J. (2008). Spatial frequency tuning of upright and inverted face identification. *Vision Research*, **48**, 2817–2826.

Goffaux, V., van Zon, J., & Schiltz, C. (2011). The horizontal tuning of face perception relies on the processing of intermediate and high spatial frequencies. *Journal of Vision*, **11**(10):1, 1–9.

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Identification of band-pass filtered letters and faces by human and ideal observers. *Vision Research*, **39**, 3537–3560.

Harmon, L. D. (1973). The recognition of faces. *Scientific American*, **229**(5), 71–82.

Hayes, T., Morrone, M. C., & Burr, D. C. (1986). Recognition of positive and negative bandpass-filtered images. *Perception*, **15**, 595–602.

Jenkins, R., Beaver, J. D., & Calder, A. J. (2006). I thought you were looking at me: direction-specific aftereffects in gaze perception. *Psychological Science*, **17**, 506–513.

Jiang, F., Blanz, V., & O'Toole, A. J. (2006). Probing the visual representation of faces with adaptation: A view from the other side of the mean. *Psychological Science*, **17**, 493–500.

Leopold, D. A., O'Toole, A. J., Vetter, T., & Blanz, V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience*, **4**, 89–94.

Loftus, G. R., & Harley, E. M. (2005). Why is it easier to identify someone close than far away? *Psychonomic Bulletin & Review*, **12**, 43–65.

Lundqvist, D., & Litton, J. E. (1998). The Averaged Karolinska Directed Emotional Faces (AKDEF). (Material developed in 1998 at Department of Clinical Neuroscience, Karolinska Institutet, Stockholm.)

Majaj, N. J., Pelli, D. G., Kurshan, P., & Palomares, M. (2002). The role of spatial frequency channels in letter identification. *Vision Research*, **42**, 1165–1184.

Näsänen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research*, **39**, 3824–3833.

Oruç, I., & Barton, J. J. (2010a). Critical frequencies in the perception of letters, faces, and novel shapes: evidence for limited scale invariance for faces. Journal of Vision, **10**(12):20, 1–12.

Oruç, I., & Barton, J. J. (2010b). A novel face aftereffect based on recognition contrast thresholds. *Vision Research*, **50**, 1845–1854.

Oruç, I., & Landy, M. S. (2009). Scale dependence and channel switching in letter identification. *Journal of Vision*, **9**(9):4, 1–19.

Parker, D. M., & Costen, N. P. (1999). One extreme or the other or perhaps the golden mean? Issues of spatial resolution in face processing. *Current Psychology*, **18**, 118–127.

Rubin, G., & Siegel, K. (1984). Recognition of low-pass faces and letters. *Investigative Ophthalmology and Visual Science*, **30**, 96.

Schyns, P. G., Bonnar, L., & Gosselin, F. (2002). Show me the features! Understanding recognition from the use of visual information. *Psychological Science*, **13**, 402–409.

Sinha, P. (2002a). Identifying perceptually significant features for recognizing faces. *Proc. SPIE: Human Vision and Electronic Imaging VII*, **4662**, 12–21, doi:10.1117/12.469529

Sinha, P. (2002b). Recognizing complex patterns. *Nature Neuroscience*, **5**, (Supplement) 1093–1097.

Smith, F. W., & Schyns, P. G. (2009). Smile through your fear and sadness: transmitting and identifying facial expression signals over a range of viewing distances. *Psychological Science*, **20**, 1202–1208.

Webster, M. A., Kaping, D., Mizokami, Y., & Duhamel, P. (2004). Adaptation to natural facial categories. *Nature*, **428**(6982), 557–561.

Willenbockel, V., Fiset, D., Chauvin, A., Blais, C., Arguin, M., Tanaka, J. W., et al. (2010). Does face inversion change spatial frequency tuning? *Journal of Experimental Psychology: Human Perception and Performance*, **36**, 122–135.

Yip, A. W., & Sinha, P. (2002). Contribution of color to face recognition. *Perception*, **31**, 995–1003.