ACCEPTED AUTHOR VERSION, TO BE PUBLISHED IN THE *QUARTERLY JOURNAL OF EXPERIMENTAL PSYCHOLOGY*

The role of the eye region for familiar face recognition: Evidence from spatial low-pass filtering and contrast negation.

Bartholomew P.A. Quinn & Holger Wiese

Durham University, United Kingdom

Author Note

Bartholomew P.A. Quinn, Department of Psychology, Durham University, United Kingdom; Holger Wiese, Department of Psychology, Durham University, United Kingdom.

Correspondence concerning this article should be addressed to Holger Wiese, Department of Psychology, Durham University, Science Site, Durham DH1 3LE, United Kingdom. E-mail: holger.wiese@durham.ac.uk.

Data is available at https://osf.io/cd4zv/. Study material consists of face photographs and can not be published due to copyright restrictions.

Abstract

What information is used for familiar face recognition? While previous research suggests a particular importance of the eye region, information from the rest of the face also needs to be integrated. What type of information is used in conjunction with the eyes is largely unclear. In three experiments, participants were asked to recognise so-called face chimeras, in which the eye region was not manipulated while the rest of the face was either presented in negative contrast (contrast chimeras) or low-pass filtered (blur chimeras). We show (i) that both chimeras are recognised substantially better than fully blurred faces, (ii) that the recognition advantage for blur chimeras is specific to the eye region but cannot be explained by cues available in this part of the face alone, and (iii) that a combination of negative contrast and blurring outside of the eye region eliminates the chimera advantage. We conclude that full-frequency but distorted surface reflectance cues (in contrast chimeras) or coarse shape information (in blur chimeras) can be used in combination with the eye region for effective face recognition. Our findings further suggest that the face recognition system can flexibly use both types of information, depending on availability.

Keywords: familiar face recognition; low-pass filtering; contrast negation; face chimeras; eye region

The role of the eye region for familiar face recognition:

Evidence from spatial low-pass filtering and contrast negation.

Humans are highly accurate at identifying familiar individuals from their faces (e.g. Wiese, Tüttenberg, et al., 2019; Young & Burton, 2017). However, despite considerable research, it is not completely understood what information is extracted from a face to accomplish this. It is well established that isolated features, such as the eyes or mouth, are not recognised effectively, and that some form of holistic or integrational processing underlies face recognition (DeGutis, Wilmer, Mercado, & Cohan, 2013; Tanaka & Farah, 1993; Young, Hellawell, & Hay, 1987). At the same time, empirical evidence suggests that the eye region is more important than other parts of the face, as masking the eyes (and eyebrows; Sadr, Jarudi, & Sinha, 2003) lowers recognition accuracy more than masking other features (e.g. McKelvie, 1976). The present study further examined the role of the eye region, as well as the integration of information from the rest of the face, for familiar face recognition by manipulating spatial frequency and contrast information, while, in the critical conditions, keeping the eye region unaffected.

Contrast negation – and particularly the reversal of luminance – substantially diminishes familiar face recognition (Galper, 1970; Kemp, Pike, White, & Musselman, 1996). While this effect was initially assumed to result from the disruption of shape-from-shading information (e.g. Johnston, Hill, & Carman, 1992; Kemp et al., 1996), more recent evidence suggests that it is primarily driven by the disruption of albedo and texture information, and therefore surface reflectance (Liu, Collin, & Chaudhuri, 2000; Russell, Sinha, Biederman, & Nederhouser, 2006). Interestingly, despite its substantial effect on face recognition, contrast negation preserves important information such as the spatial frequency spectrum of the original image, and its deleterious effect has been interpreted as reflecting the inhibition rather than elimination of recognition cues (see Sormaz, Andrews, & Young, 2013, for a related discussion).

Of particular relevance for the present study, contrast negation disrupts the ordinal contrast relationship of faces, with normally darker, sunken eye sockets, relative to the surrounding brighter face regions (i.e. cheeks, forehead). This contrast relationship is stable across naturally occurring lighting conditions and viewpoints (Braje, Legge, & Kersten, 2000), and has been suggested as an important cue for face recognition. In line with this idea, recognition accuracy is negatively affected when ordinal contrast is diminished by lighting a face from below (e.g. Johnston et al., 1992). This deficit is largely restored when a below-lit face is additionally contrast inverted, which re-establishes the darker eye region (Liu, Collin, Burton, & Chaudhuri, 1999).

Similarly, contrast chimeras – stimuli with positive eye regions in an otherwise negative face (see Figure 1) – are substantially easier to recognise than their fully negative counterparts (Fisher, Towler, & Eimer, 2016; Gilad, Meng, & Sinha, 2009; Sormaz et al., 2013; see also Wiese, Chan, & Tüttenberg, 2019). Again, this advantage has been explained by the restoration of the ordinal contrast relationship by contrast chimeras, which is disrupted by contrast negation (Gilad et al., 2009). In line with this suggestion, rendering other features in positive contrast results in a markedly smaller advantage (Sormaz et al., 2013), which might also be interpreted as reflecting a specific salience of the eye region for effective recognition. Critically, the advantage for eye chimeras over negative faces does not arise from the availability of recognition cues within the eye region alone, as isolated eyes or eyes in dark silhouettes are recognized substantially more poorly than contrast chimeras (Gilad et al., 2009; Sormaz et al., 2013).

While ordinal contrast information has been identified as an important cue to face recognition, manipulations that leave this relationship intact can nonetheless disrupt face recognition. Of particular relevance, removing high- (HSF) and mid-spatial frequencies (MSF) through low-pass filtering (blurring) limits an image to only coarse low-spatial frequency (LSF) information. Face recognition has been shown to operate preferentially at 8-16 cycles-per-image (cpi; see Ruiz-Soler & Beltran, 2006, for a review). Consequently, low-pass filtering faces below an 8cpi cut-off degrades recognition accuracy incrementally (Costen, Parker, & Craw, 1996; Nasanen, 1999; Parker & Costen, 1999), while keeping the broad ordinal contrast information of the face largely intact.

In sum, blurring and contrast negation affect face recognition differently. On the one hand, detailed (i.e., HSF and MSF) cues are removed from a face by blurring but preserved by contrast negation. On the other hand, broad ordinal contrast relationships are disrupted by contrast inversion, while blurring retains the overall (albeit coarser) patterns of shading of an unfiltered face. Previous research on contrast chimeras (Gilad et al., 2009; Sormaz et al., 2013) has concluded that the advantage of these stimuli over full negative faces is based on the restoration of typical ordinal contrast relationships. If this restoration was the primary factor driving the observed chimera advantage, an unfiltered eye region in a blurred face (see Figure 1) should not substantially improve recognition relative to fully blurred faces, as both fully filtered faces and such blur chimeras retain typical ordinal contrast relationships. By contrast, establishing an advantage for blur chimeras over fully filtered faces would suggest that a chimera effect can occur even when the unmanipulated eye region does not restore ordinal contrast. As detailed below, blur chimeras also allow further examining the contribution of the eye region relative to the rest of the face in familiar face recognition.

The present study examined face recognition with contrast and blur chimeras in a series of three experiments. Experiment 1 examined a relatively large group of participants to establish the basic chimera effect for spatially low-pass filtered faces. Experiment 2 then directly compared the recognition of blur and contrast chimeras while at the same time investigating the specificity of the effect to the eye region. Finally, Experiment 3 examined face recognition with combined blur and contrast manipulations to test whether a chimera effect can be observed in these severely degraded stimuli, and whether the re-establishment of ordinal contrast would be helpful even for low-pass filtered stimuli.

Experiment 1

Experiment 1 was designed to establish initial evidence of a potential advantage of blur chimeras over fully low-pass filtered faces. We reasoned that if the advantage for eye chimera stimuli was exclusively dependent on the re-establishment of ordinal contrast relationships, then no advantage for blur chimeras relative to fully filtered images should be observed. Alternatively, if an unmanipulated eye region added to the information preserved by low-pass filtering (such as coarse 3D shape cues), then blur chimeras should be recognised more accurately than fully blurred faces.

Method

Participants

Prior to data collection, the required sample size was calculated using G*Power (Faul, Erdfelder, Lang, & Buchner, 2007), assuming a medium effect size for an independent-samples *t*-test (difference between blur faces and chimera faces, Cohen's $d = 0.5$, power = .80, two-tailed $\alpha = .05$). This power analysis revealed a minimum sample size of 64 participants per group. The effect size was assumed to be smaller than that generated by contrast chimeras (e.g. Sormaz et al., 2013) due to the less detrimental effect of moderate low-pass filtering than negative contrast on recognition accuracy (Sandford, Sarker, &

Bernier, 2018). A total of 191 Durham University undergraduate students were tested, 15 of which were excluded due to insufficient familiarity with the presented celebrity faces (less than 20 [i.e. 50%] correct identifications in the unfiltered face condition, see below), or for not having followed the instructions outlined in the experiment. 176 participants (153 female, mean age = 19.4 years, *SD* = 1.2) were retained in the final sample. Participants were divided pseudo-randomly into two groups, each consisting of 88 participants, resulting in an achieved power of .99 with the above parameter values. All participants provided written informed consent and were compensated with participant pool credit. The experiment was approved by the ethics committee of Durham University's Psychology department.

Stimuli

84 images, comprising of two different images of 42 celebrities (i.e. actors, politicians, musicians) were collected using Google Image search. Images were standardized using Adobe Photoshop CS6 (Version 13.0.1; www.adobe.com). Faces were cropped from backgrounds, converted to greyscale, and pasted onto a uniform grey background. Cropped faces were scaled to a standardized size of 228x342 pixels.

One image of each of the 42 celebrities was not manipulated further ('Unfiltered faces'). The second image was spatially low-pass filtered using FourierImage (2017 version; www.nasanen.info; low-pass Butterworth Filter, filter exponent: 5, Cut-off frequency: 7cpi) to create 'Blur faces'. In addition, 'Blur chimeras' of each celebrity's second image were created by fitting lemniscates with smoothed edges around the eyes and eyebrows of each face (Adobe Photoshop CS6; Refine Edge Tool, Smooth: 70, Feather: 4.5px, Shift Edge: 12%) and leaving the selected region unfiltered (akin to Gilad et al., 2009). Exemplars of the three image conditions are depicted in Figure 1.

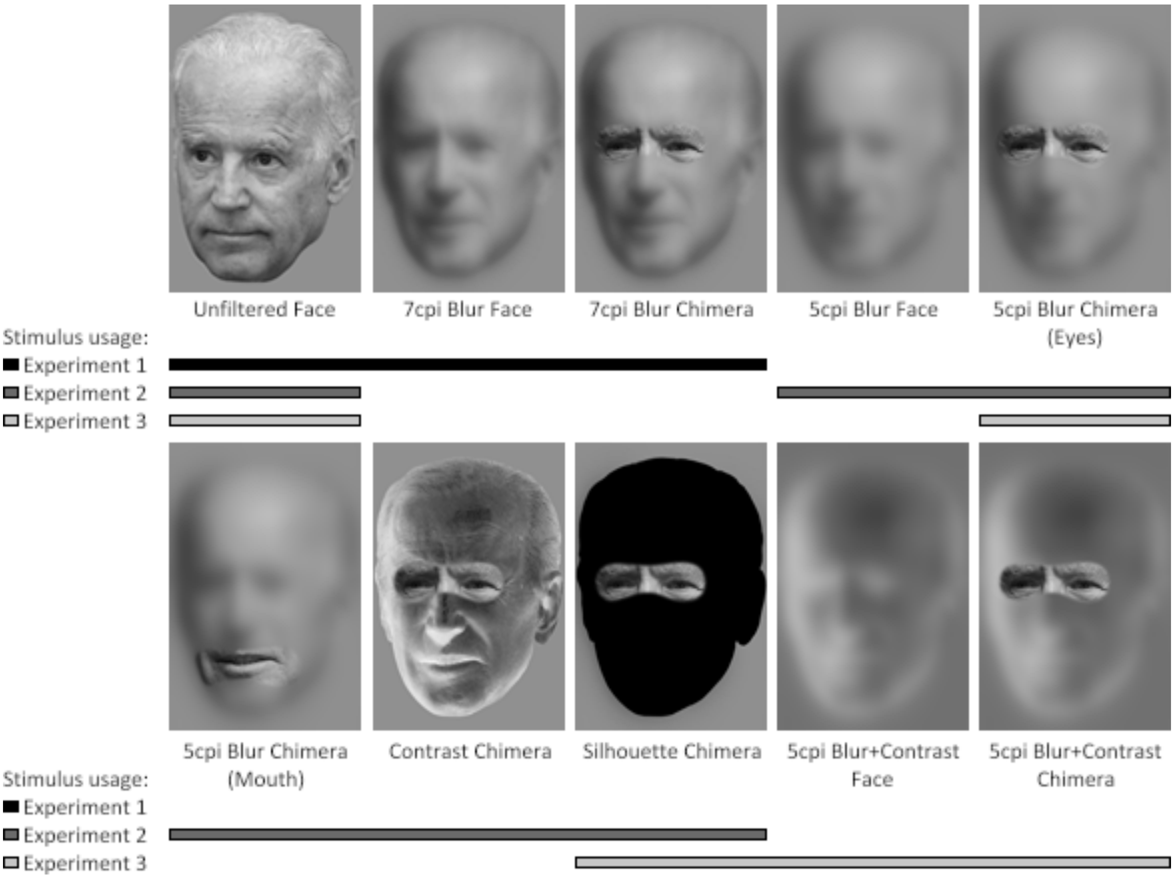

Figure 1. Examples an original image and the nine manipulations used in Experiment 1-3 for one celebrity (U.S. President Joe Biden), with indicators of stimulus usage across experiments. Note that a different image was used for the unfiltered face stimulus of each identity.

Procedure

The experiment was programmed in PsychoPy (Version 2020.2.10; www.psychopy.org), and presented online via Pavlovia (www.pavlovia.org), allowing participants to take part using their personal computers. Participants were pseudo-randomly allocated to one of two image manipulation conditions, either featuring blur faces or blur chimeras.

The experiment began with a practice block featuring four faces from two celebrity identities not used elsewhere in the experiment (two unfiltered faces, and two in the participant's relevant image manipulation condition) to familiarize them with the task. For the first block of testing, participants then either viewed 40 trials featuring blur faces or blur chimeras, depending on their group, with each celebrity presented once in random order. Trials started with a white fixation cross in the centre of the screen (500ms), followed by a face in the same location. 500ms after stimulus onset, a text box and prompt appeared while the face stimulus remained on the screen, requesting participants to indicate if they recognised the depicted individual. If participants did recognize the individual, they were asked to type in the name, stage name, or relevant identity-specific information (e.g. a specific role they had played in a film). Else, they were prompted to press the 'up arrow' key to move on to the next trial. All participants then completed a second block featuring the previously unseen images of the same 40 celebrities in the unfiltered face condition presented in randomized order. Participants were again requested to identify each of the faces using the text box.

Participants were assumed to have recognized a face if they provided an accurate name or identity-specific information about the depicted individual (e.g. "the actor who plays James Bond" rather than "actor", "British prime minister" rather than "politician"). The proportion of correct target identification was calculated by dividing the number of correct responses in the blur face/chimera condition by the number of correct responses in the unfiltered face condition. In a small number of instances, participants recognized a face in the blur face/chimera condition, but not in the unfiltered face condition. These were treated as recognition across both conditions[1]. For statistical analysis, an independent samples *t*-test was

[1] We have kept these trials (0.36%, 0.5%, and 0.29% in Experiments 1-3, respectively) in the analysis as a correct identification in the more difficult (e.g. fully blurred or blur chimera) conditions indicated to us that the participant was in fact familiar with the identity. Failure to identify the face in the unmanipulated condition (e.g. due to atypical images or erroneous key presses) would then be irrelevant. However, exclusion of these trials does not change the outcome of any of the reported statistical tests.

conducted. Confidence intervals and bias-corrected effect size measures, $d_{unb.}$ (Cumming, 2012), are reported.

The study design, hypotheses, and analysis plan for the experiments presented here were not preregistered. All data is publicly available on the Open Science Framework website (https://osf.io/cd4zv/). Celebrity faces are not made publicly available for copyright reasons, but examples (licensed under the Creative Commons Public Domain Mark 1.0 license) are given in Figure 1.

Results

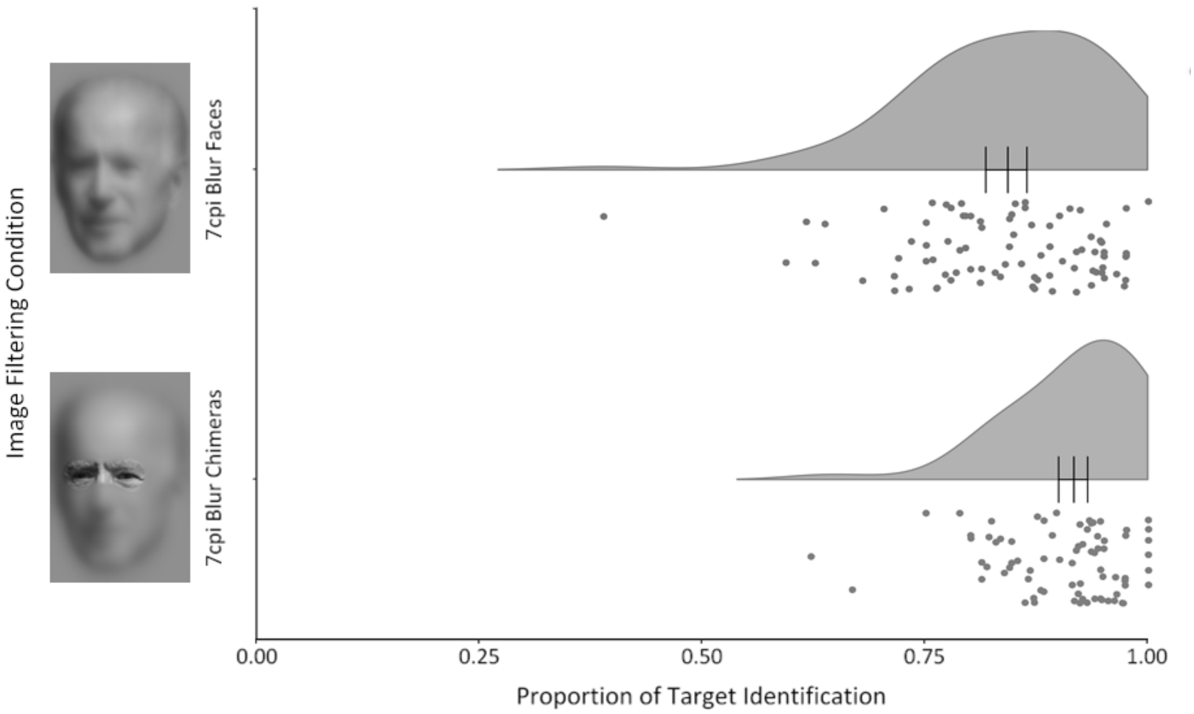

Figure 2. Mean (vertical lines) and individual (dots) proportions of correct target identification for the Blur Face and Chimera Face conditions. Error bars reflect 95% confidence intervals.

Results are depicted in Figure 2. Levene's test demonstrated that variances between Blur Faces and Blur Chimeras were not equal, $F = 9.16$, $p = .003$. Accordingly, a t statistic not assuming homogeneity of variance was calculated. The independent samples t-test indicated that participants in the Blur Chimeras condition ($M = 0.92$, $SD = 0.08$) had a significantly higher proportion of correct target identifications than participants in the Blur Faces condition ($M = 0.84$, $SD = 0.11$), $t(156.24) = 5.31$, $p < .001$, $d_{unb} = .80$, 95% CI [0.50, 1.12].

Discussion

Experiment 1 demonstrates that an unfiltered eye region in an otherwise blurred face provides a significant enhancement in recognizability relative to fully blurred faces. This finding shows that a chimera advantage is not limited to contrast negation but extends to a manipulation which on the one hand reduces recognizability but on the other hand leaves ordinal contrast relationships intact. This advantage of blur chimeras over fully low-pass filtered faces therefore cannot be explained by the restoration of such contrast relationships.

The results also provide initial information regarding the usage of information from the filtered and unmanipulated face regions. Blurring eliminates recognition cues based on spatial detail (Loftus & Harley, 2005; Vuilleumier, Armony, Driver, & Dolan, 2003), however retains others, such as broad contrast information indicative of coarse 3D shape (A. Hayes, 1988; T. Hayes, Morrone, & Burr, 1986), and above-chance recognition in the fully filtered condition is likely based on such cues. Importantly, the advantage for blur chimeras over fully filtered images implies that the availability of HSF and MSF information in the eye region substantially improves recognition.

However, before any more substantive conclusions about the nature of the observed effect can be drawn, a number of open questions need to be addressed. First, previous studies have indicated that isolated eyes, or positive eyes on dark silhouettes, are poorly recognized when compared to unmanipulated faces (e.g. Sormaz et al., 2013), but their recognition rate had not been established in comparison to the current experimental manipulations. It is therefore not known whether the chimera advantage in Experiment 1 was driven by the unfiltered eye region alone. Second, from Experiment 1 alone, the specificity of the observed effect to the eye region is unclear. The observation of no recognition advantage from an unfiltered mouth-region in an otherwise blurred face would add further evidence to the suggestion that the eye region is particularly important for familiar face recognition. Finally, it would be beneficial to examine contrast and blur chimeras within the same experiment to test the extent to which the two manipulations are comparable. Experiment 2 was designed to resolve these questions.

Experiment 2

In Experiment 2, we again examined the recognition of blur chimeras, as well as unfiltered eyes in dark silhouettes (silhouette chimeras) and unfiltered mouth regions in otherwise blurred faces (mouth chimeras). These latter stimuli allowed us to estimate the contribution of the low-pass filtered face parts in blur chimera recognition, as well as the specificity of the effect to the eye region, respectively. We further directly compared contrast and blur chimeras. As contrast negation has a more drastic effect on face recognition relative to moderate low-pass filtering (Sandford et al., 2018), we reasoned that the application of a more severe low-pass filter (with a 5cpi instead of 7cpi cut-off) would give more comparable results to contrast negation.

Assuming a contribution of the low-pass filtered face region to face recognition even in these more drastically blurred images, we predicted a clear recognition advantage for blur chimeras over silhouette chimeras. Moreover, assuming that the chimera advantage is specific to the eye region, we predicted no advantage for mouth chimeras over fully blurred faces. Finally, assuming that contrast negation and low-pass filtering with a 5cpi cut-off would produce stimuli that are similarly difficult to recognise, and given that the unmanipulated eye region was identical in both stimulus categories, we did not expect any advantage of contrast over blur chimeras.

Method

Participants

The required sample size was estimated based on the effect size observed in Experiment 1 using G*Power (Faul et al., 2007), assuming a large effect size for the difference between any two of the tested conditions (independent samples *t*-test, $d = 0.8$, power = .80, two-tailed $\alpha = .05$). This revealed a required sample size of 24 participants per group. A total of 163 undergraduate students at Durham University were tested, 33 of which were excluded due to insufficient familiarity with the presented celebrities faces (less then 20 correct identifications in the unfiltered face condition). 130 participants (116 female, mean age = 19.6 years, $SD = 1.1$), 26 per group, were retained in the final sample. All participants provided written informed consent and were compensated with participant pool credit. The experiment was approved by the ethics committee of Durham University's Psychology department.

Stimuli

84 images, comprising of two different images of 42 celebrities not used in Experiment 1 were collected and standardized as described above. As before, one image per celebrity was not edited and used in the 'Unfiltered faces' condition. The second image was used for the five image manipulation conditions. 'Blur faces' and 'Blur chimeras (eyes)' were created analogously to Experiment 1, with the exception of using a filter cut-off frequency of

5cpi instead of 7cpi. ‘Blur chimeras (mouth)’ were created using the 5cpi blur chimera (eyes) images as a template and moving the unfiltered lemniscate from the eye region to the mouth region of each face. This procedure was applied individually to each stimulus, so that mouth chimeras retained the same number of unfiltered pixels as eye chimeras. ‘Contrast chimeras’ and ‘Silhouette chimeras’ were created by replacing the 5cpi low-pass filtered image behind the unfiltered eye layer with an unfiltered face rendered in negative contrast or 0% luminance respectively. Exemplars are depicted in Figure 1.

Procedure

The experiment was conducted analogously to Experiment 1, the only exception being that for each group, one of the five novel image manipulation conditions was used in block 1.

Results

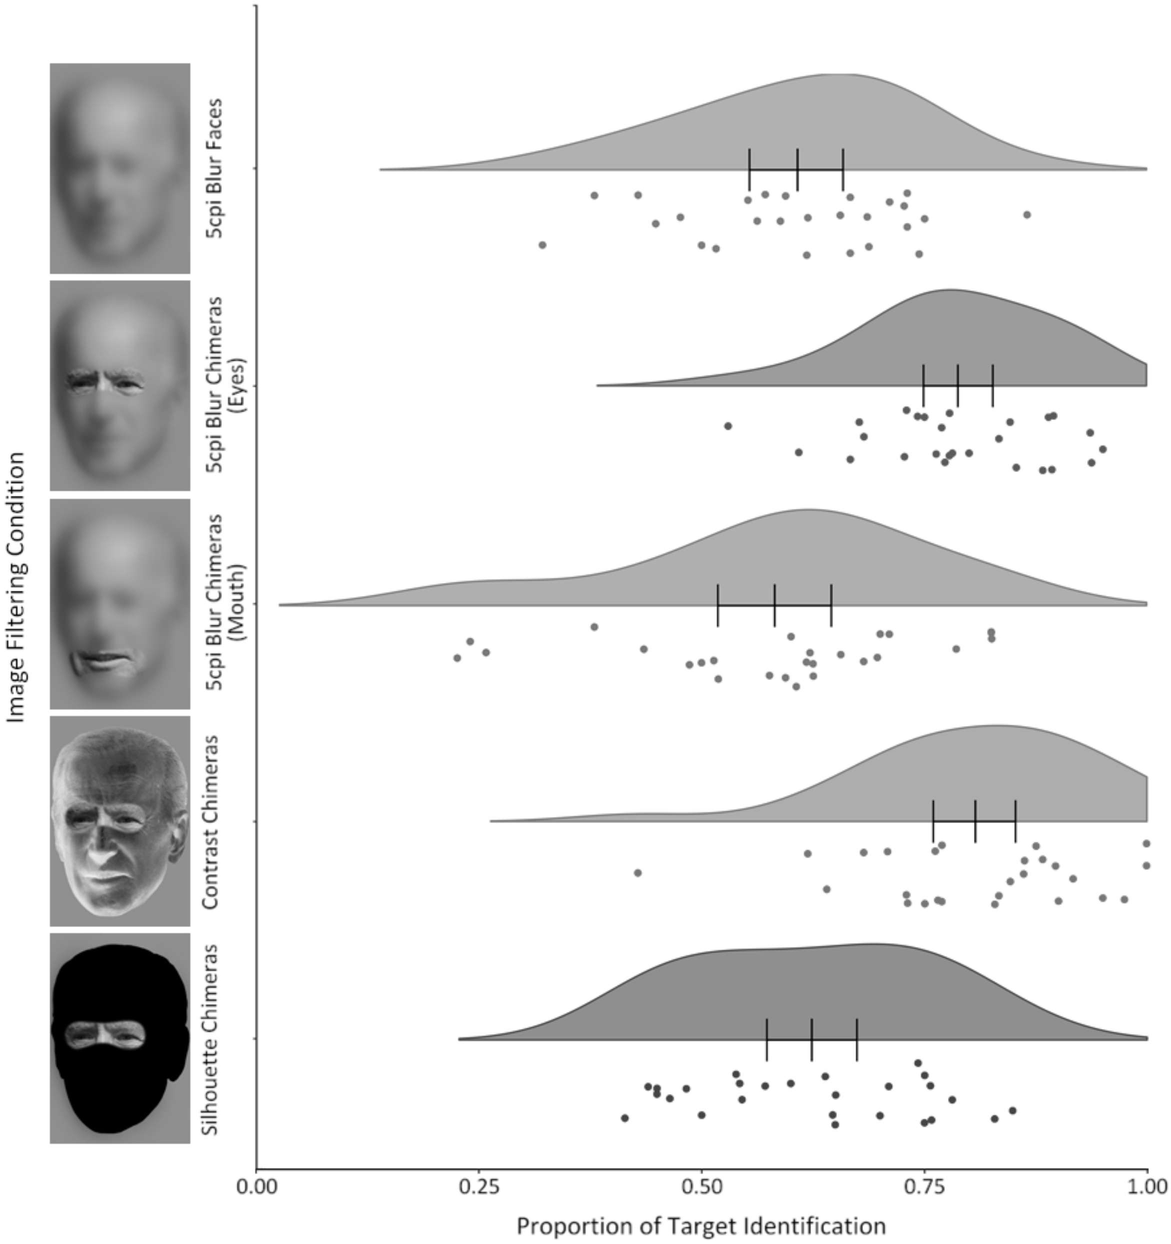

Figure 3. Mean (vertical lines) and individual (dots) proportions of correct target identification for five image manipulation conditions. Error bars reflect 95% confidence intervals.

Results are depicted in Figure 3. A one-way independent-measures ANOVA revealed a significant difference in the proportion of correct target identifications between the five image manipulation conditions, $F(4,125) = 16.30$, $p < .001$, $\eta^2 = .37$. Follow-up comparisons

(Table 1) yielded no significant difference between blur chimeras (eyes) (*M* = 0.79, *SD* = 0.10) and contrast chimeras (*M* = 0.81, *SD* = 0.13). Both of these groups had significantly higher proportions of correct target identifications than participants in the blur chimeras (mouth) condition (*M* = 0.58, *SD* = 0.17), the silhouette chimeras condition (*M* = 0.62, *SD* = 0.13), and the blur faces condition (*M* = 0.61, *SD* = 0.13), between which there was also no significant difference.

As one prediction for Experiment 2 was that blur and contrast chimeras would not differ in their recognition rates, an additional Bayesian independent-samples *t*-test was conducted. This test yielded moderate evidence for the null hypothesis, BF_{01} = 3.01 (error = 0.02%). A second prediction was that blurred faces and silhouette chimeras would also not differ in their recognition rates. An additional Bayesian independent-samples *t*-test was conducted. This test yielded moderate evidence for the null hypothesis, BF_{01} = 3.32 (error = 0.02%).

Table 1

Independent-samples comparisons of the effects of image filtering condition on proportion of correct target identification scores in Experiment 2.

Effect	M_{diff}	95% CI	*t*(50)	*p*	d_{unb}	95% CI
Blur faces vs. blur chimeras (eyes)	0.18	[0.11, 0.25]	5.52	**<.001**	1.51	[0.94, 2.21]
Blur faces vs. blur chimeras (mouth)	-0.03	[-0.11, 0.06]	-0.61	.543	-0.17	[-0.73, 0.38]
Blur faces vs. contrast chimeras	0.20	[0.13, 0.27]	5.54	**<.001**	1.51	[0.94, 2.22]
Blur faces vs. silhouette chimeras	0.02	[-0.06, 0.09]	0.44	.661	0.12	[-0.43, 0.68]
Blur chimeras (eyes) vs. blur chimeras (mouth)	-0.21	[-0.28, -0.13]	-5.26	**<.001**	-1.44	[-2.13, -0.87]
Blur chimeras (eyes) vs. contrast chimeras	0.02	[-0.05, -0.09]	0.60	.550	0.16	[-0.39, 0.73]
Blur chimeras (eyes) vs. silhouette chimeras	-0.16	[-0.23, -0.10]	-4.96	**<.001**	-1.36	[-2.04, -0.79]
Blur chimeras (mouth) vs. contrast chimeras	0.23	[0.14, 0.31]	5.36	**<.001**	1.46	[0.90, 2.16]
Blur chimeras (mouth) vs. silhouette chimeras	0.04	[-0.04, 0.13]	0.99	.329	0.27	[-0.28, 0.84]

Contrast chimeras vs. silhouette chimeras	-0.18	[-0.26, -0.11]	-5.04	**<.001**	-1.38	[-2.06, -0.81]

Discussion

Experiment 2 found that those participants who viewed blur or contrast chimeras with unmanipulated eyes were considerably more accurate than participants in any other condition, which replicates our and others' previous results. Critically, both silhouette and mouth chimeras were recognised substantially worse than eye chimeras. Finally, recognition accuracy between contrast and blur (eyes) chimeras did not differ, suggesting comparable usability of facial information with a 5cpi low-pass filter in the latter condition.

Experiment 2 was designed to resolve the questions left open by our first experiment. First, performance in the blur chimera condition cannot be explained by the unmanipulated eye region alone, as presenting unfiltered eyes in dark silhouettes did not result in comparable recognition rates. Consequently, blur chimera recognition reflects a combination of cues from both filtered and unfiltered face regions. Second, as mouth chimeras were not recognized better than fully blurred faces, the increased recognition of eye chimeras critically depends on the eye region. In other words, HSF/MSF information from the mouth region does not contribute relevant cues over and above those available in low-pass filtered images. Finally, similar performance levels for contrast relative to blur chimeras suggest comparable usage of recognition cues both from the eye region (which was identical in the two conditions) and from the manipulated part of the stimuli. This finding indicates that the visual system seems to flexibly use those recognition cues that are available in the manipulated face part (such as full frequency spectrum but distorted surface reflectance information in contrast chimeras versus only LSF information in blur chimeras). It also further strengthens our observation that the reestablishment of ordinal contrast is not a necessary prerequisite for the occurrence of an (eye) chimera advantage.

Experiment 3 was designed to further examine the type of information used for the recognition of the different chimera stimuli. As outlined above, full frequency spectrum but distorted surface reflectance information seems to be extracted from the manipulated part of contrast chimeras, while only spatially coarse LSF information can be used for blur chimeras. However, particularly given the similar recognition rates in Experiment 2, one might alternatively assume that the same information is extracted from the manipulated regions in both types of chimeras. This would suggest, however, that whatever cues are extracted are neither affected by low-pass filtering nor contrast negation, and should therefore still be available in faces that are both blurred and presented in negative contrast. Experiment 3 examined this alternative explanation by testing recognition of combined blur and contrast chimeras.

In addition, we reasoned that combined blur and contrast chimeras could assist in further examining the role of ordinal contrast information for face recognition. The combined application of low-pass filtering and contrast negation to a face would disrupt ordinal contrast (as it includes the reversal of luminance). However, presenting such a face with an unmanipulated eye region, i.e. a combined blur and contrast chimera, would re-establish this contrast relationship. If a recognition advantage for such stimuli was observed, the effect of ordinal contrast would likely be contingent on broad LSF information. However, if a recognition advantage for combined chimeras was absent, the contribution of ordinal contrast to face recognition would depend on the availability of HSF/MSF information.

Experiment 3

To examine the research questions outlined above, Experiment 3 tested the recognition of combined blur and contrast chimeras as well as fully filtered contrast negative faces. To estimate the effect of negative contrast on the recognition of blur chimeras and allow testing for contributions of the manipulated face region, 5cpi blur chimeras and silhouette chimeras from Experiment 2 were added to the analysis. Assuming that different cues were extracted from the manipulated parts of contrast and blur chimeras, and that both types of cues would be effectively eliminated by the combination of image manipulations, we expected no recognition advantage of combined blur and contrast chimeras over the silhouette condition (representing recognition from the eye region alone).

Method

Participants

Based on the power analysis for Experiment 2, a total of 57 additional participants were tested. Of these, 5 were excluded due to insufficient familiarity with the celebrities (less than 20 correct identifications in the unmanipulated condition). 52 participants (43 female, mean age = 20.2 years, *SD* = 1.3), 26 per group, were retained in the final sample. All participants provided written informed consent and were compensated with participant pool credit. The experiment was approved by the ethics committee of Durham University's Psychology department.

Stimuli

Experiment 3 used the same stimulus set as Experiment 2. To create blur+contrast faces, images from the 5cpi blur face condition in Experiment 2 were contrast inverted. The same unfiltered eye lemniscates used in Experiment 2 were layered on top of these images to create blur+contrast chimeras. Examples are depicted in Figure 1.

Procedure

Participants either saw blur+contrast faces or blur+contrast chimeras in block 1. Other than this change, the experiment was conducted analogously to Experiments 1 and 2. In addition to direct comparisons of the two conditions introduced in Experiment 3, blur+contrast faces and chimeras were compared to blur chimera (eyes) and silhouette chimera conditions from Experiment 2.

Results

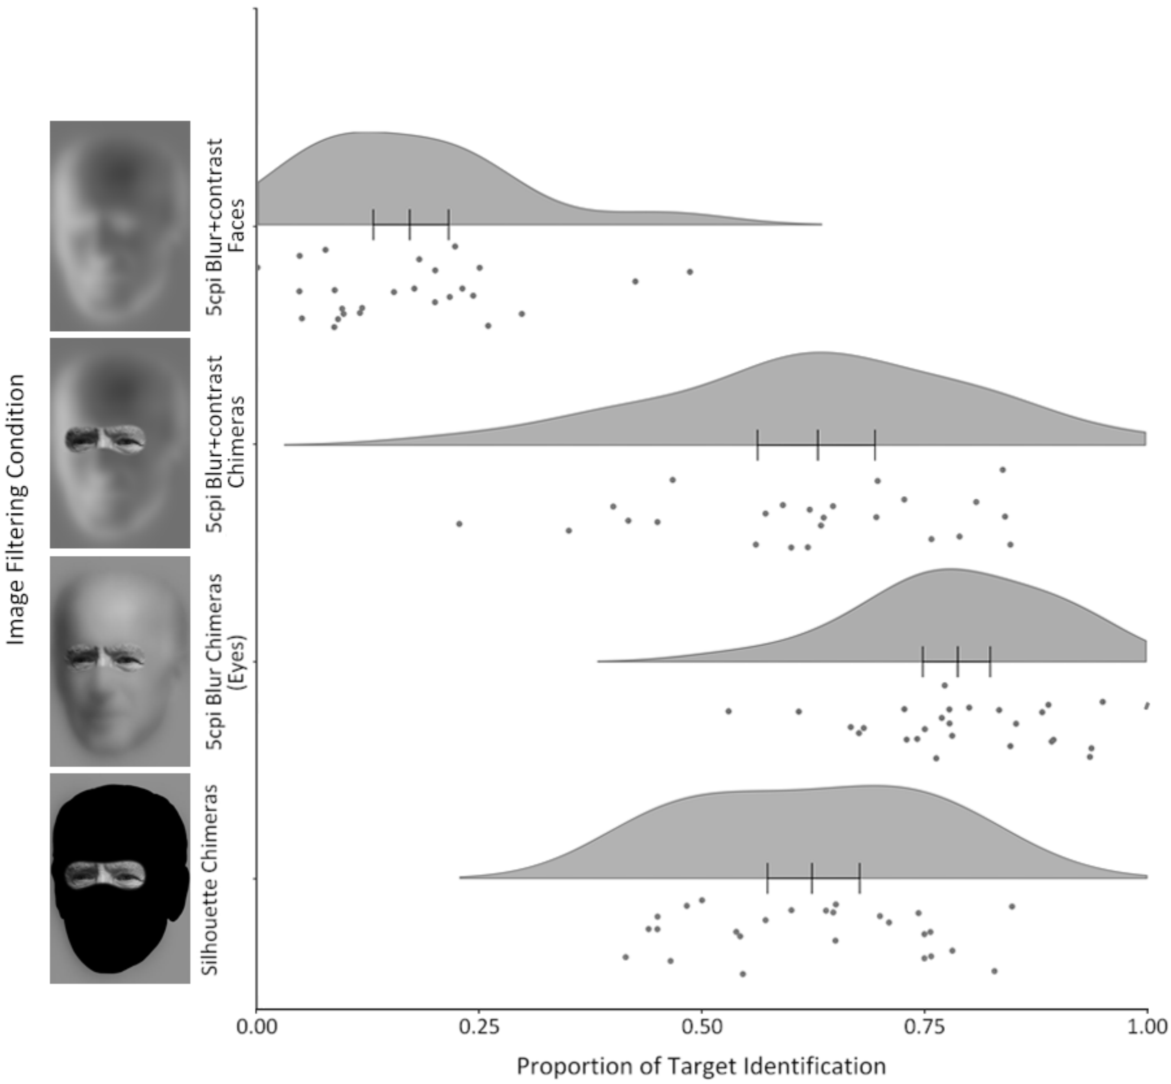

Figure 4. Mean (vertical lines) and individual (dots) proportions of correct target identification for five image manipulation conditions. Error bars reflect 95% confidence intervals.

Results are depicted in Figure 4. A one-way independent-measures ANOVA revealed a significant effect, $F(3,100) = 101.17$, $p < .001$, $\eta^2 = .75$. Follow-up comparisons (see Table 3) indicated that participants in the blur+contrast face group (*M* = 0.17 , *SD* = 0.12) had significantly lower proportions of correct target identifications than those in the blur+contrast chimera (*M* = 0.63, *SD* = 0.18), blur chimera (eyes) (*M* = 0.79, *SD* = 0.10), and silhouette chimera (*M* = 0.62, *SD* = 0.13) groups. Proportions of correct target identifications in the blur+contrast chimera group were significantly lower than scores from participants in the blur chimera (eyes) group, while, critically, revealing no difference from the scores of participants in the silhouette chimera condition.

To more directly test the prediction of no difference between blur + contrast chimeras and silhouette chimeras (against the alternative hypothesis of better recognition in the former condition), an additional Bayesian independent-samples *t*-test was conducted, which revealed moderate evidence for the null hypothesis, $BF_{0+} = 3.21$ (error = 0.01%).

Table 3

Independent-samples comparisons of the effects of image filtering condition on proportion of target identification scores.

Effect	M_{diff}	95% CI	*t*(50)	*p*	d_{unb}	95% CI
Blur+contrast faces vs. blur+contrast chimeras	0.46	[0.38, 0.54]	11.13	**<.001**	3.04	[2.34, 4.00]
Blur+contrast faces vs. blur chimeras (eyes)	0.62	[0.56, 0.68]	20.25	**<.001**	5.53	[4.50, 7.01]
Blur+contrast faces vs. silhouette chimeras	0.45	[0.38, 0.52]	13.17	**<.001**	3.60	[2.83, 4.66]
Blur+contrast chimeras vs. blur chimeras (eyes)	0.16	[0.08, 0.24]	3.92	**<.001**	1.07	[0.52, 1.72]
Blur+contrast chimeras vs. silhouette chimeras	-0.01	[-0.09, 0.80]	-0.15	.88	0.04	[-0.60, 0.51]

Experiment 3 Discussion

While combined blur and contrast chimeras were recognized more accurately than fully blurred contrast-negative faces, they yielded no advantage over silhouette chimeras. This finding suggests that a combination of blurring and contrast negation severely disrupted the information available in the manipulated parts of simple contrast and blur chimeras. It therefore appears that the face recognition system uses different information from the manipulated part of the image to increase recognition performance in simple contrast and blur chimeras.

We note that it remains possible that the combination of the two image manipulations had an additive effect on the same singular cue, decreasing it to a level at which it was no longer usable. However, given the different means by which blurring and contrast negation affect recognition cues, this possibility seems unlikely, and it appears more plausible that full-spectrum (including detailed HSF/MSF cues) but distorted surface reflectance information is extracted from the negative parts of contrast chimeras, whereas only coarse spatial information is used in the case of blur chimeras.

The results of Experiment 3 further suggest that the positive effect of re-establishing ordinal contrast depends on the availability of HSF/MSF information. While combined chimeras were recognised more accurately than fully filtered contrast negative faces, this difference is best explained by recognition cues provided by the unmanipulated eye region itself, rather than any re-established relation between different regions, as combined chimeras were not recognised better than silhouette chimeras. It thus appears that the combination of low-pass filtering and contrast negation effectively eliminated most face recognition cues, and that the restoration of ordinal contrast was not beneficial in this situation.

General Discussion

What information is used for familiar face recognition? The present series of experiments complements previous research by suggesting a particularly prominent role of the eye region relative to other parts of the face. Importantly, however, as the eye region alone is not sufficient for accurate recognition, some information from the rest of the face needs to be integrated, and our findings provide evidence on what type of additional information can be used. In three consecutive experiments, we tested face recognition using blur and contrast chimeras, and demonstrate that neither the availability of HSF/MSF information in the manipulated parts of the image nor the restoration of ordinal contrast information per se are essential. Instead, face recognition appears to be able to flexibly use either full-spectrum but distorted surface reflectance information or more coarse spatial cues, depending on availability. These findings are discussed in more detail below.

In line with previous studies (Gilad et al., 2009; McKelvie, 1976; Sadr et al., 2003; Sormaz et al., 2013), our experiments demonstrate the particular relevance of the eye region for familiar face recognition. This conclusion is based on the clear advantage for blur chimeras that leave the eye region unmanipulated over fully filtered faces, while at the same time an unfiltered mouth region does not elicit a comparable effect. Previous studies have indicated that the mouth contains some cues relevant for familiar face recognition (e.g. McKelvie, 1976), but these cues do not seem to be sufficient to allow an advantage over fully blurred images. It is possible that recognition cues in the mouth region are mostly based on low spatial frequencies, which are available in both mouth chimeras and fully filtered faces. Alternatively, familiar face representations could be centred around the eye region, and unmanipulated eyes, but not other features, could allow the coding and integration of additional information from the rest of the face (see Bruce & Young, 2012, p.269, for a related discussion).

Crucially for this interpretation, and in line with previous work on contrast chimeras (Gilad et al., 2009; Sormaz et al., 2013), the recognition advantage for blur chimeras is not exclusively based on the availability of fine detail in the eye region itself. Instead, the substantially more accurate recognition of blur chimeras relative to eyes presented in dark silhouettes demonstrates that information from the manipulated part of the face is also used in the former condition. Previous studies have suggested that the holistic integration of information from the eye region and other parts of the face underlies the accurate recognition of contrast chimeras (Sormaz et al., 2013). While a similar holistic mechanism also appears plausible for blur chimeras, the present experiments do not allow us to draw any such conclusions. More specifically, our results cannot differentiate between a simple additive effect of information usage from the filtered and unfiltered regions of the face, as compared to the super-additive integration assumed for holistic processing (in which the whole is more than the sum of its parts). In any case, the present results show that filter manipulations are unlikely to effectively isolate information from specific facial features, which has been suggested in recent ERP work (Mohr, Wang, & Engell, 2018).

The present results also yielded insight into the type of information underlying chimera advantages. Extending previous explanations of the effect (Gilad et al., 2009), our results indicate that the re-establishment of ordinal contrast relationships is not necessary to observe an advantage for eye chimeras relative to fully manipulated faces. The broad contrast relationship of lower luminance in the eye region compared to its surroundings is clearly important to recognizing faces (e.g. Johnston et al., 1992). However, ordinal contrast is maintained both in fully blurred faces and blur chimeras. The advantage of the latter over the former condition therefore cannot be explained by its re-establishment.

Relatedly, previous research has suggested that detailed contrast information and pigmentation are used as the primary cues for familiar face recognition (Bruce & Langton, 1994; Kaufmann & Schweinberger, 2012; Russell et al., 2006). It therefore appears plausible that such detailed (i.e. HSF/MSF) information is integrated with the eye region to recognise contrast chimeras. Blurring, however, substantially reduces or even eliminates HSF/MSF cues. The comparable recognition advantage between 5cpi blur chimeras and contrast chimeras (as observed in Experiment 2) therefore indicates that access to HSF/MSF information is not a necessary prerequisite for a chimera effect. Moreover, the similarity of the two conditions at least partly reflects comparable usage of recognition cues in the manipulated part of the face, given that (i) the unmanipulated eye region in both chimera conditions was identical, and (ii) that chimera recognition is partly based on information beyond the eye region. Therefore, cues based on coarse LSF information, which have been shown to provide a recognisable approximation of 3D shape for objects (A. Hayes, 1988; T. Hayes et al., 1986), are presumably extracted from the manipulated parts of blur chimeras. While research has emphasised the importance of surface reflectance for familiar face recognition (Bruce & Langton, 1994; Russell & Sinha, 2007), 3D shape cues are also known to contribute (e.g. Bruce, Doyle, Dench, & Burton, 1991; O'Toole, Vetter, & Blanz, 1999). Our results therefore suggest that detailed information from the eye region can be used in combination with coarse 3D shape cues from the filtered face region to allow for effective recognition.

Finally, our results suggest that different cues can be used in different circumstances. This idea was tested in our third experiment. We assumed that if the recognition advantage of blur and contrast chimeras relied on identical cues in the manipulated face region, then these cues should be unaffected by both contrast negation and low-pass filtering. Accordingly, if the two manipulations were combined, a chimera advantage should still be observed. By contrast, however, combined blur and contrast chimeras were not recognized better than silhouette chimeras. This suggests that any advantage over fully manipulated images was most probably related to the eye region itself rather than either cues from the manipulated face region or an integration of information across different face regions. Accordingly, different and complementary cues seem to be flexibly extracted from a face stimulus depending on availability.

A potential limitation of our experiments stems from their presentation in an online format, which became necessary due to the COVID 19 pandemic and the resulting impossibility to test in a controlled lab environment. Consequently, our participants used different devices to run the experiments, and no experimental control over presentation parameters such as visual angle or effective brightness of the stimuli was possible. While the extent to which our results are affected by this lack of control is unclear, we note that we tested a relatively large sample in Experiment 1, which might compensate for potential additional noise in the data introduced by the online procedure. Moreover, across the three experiments we observed strong and unambiguous effects, suggesting that their occurrence does not rely on the strict experimental control of presentation parameters. Nevertheless, we acknowledge that replication of the present finding in more controlled conditions is desirable.

In conclusion, the present study provides evidence for the particular relevance of the eye region for familiar face recognition. However, cues from the eye region alone are insufficient for effective recognition and need to be integrated with other parts of the face. Importantly, different types of recognition cues can be used in different circumstances, depending on their availability. More specifically, the manipulations tested here suggest that both coarse spatial information informing about 3D shape and detailed HSF/MSF cues can be used, and that these different cues can be similarly effective. The present results therefore contribute to our understanding of familiar face recognition by demonstrating the flexibility of the involved processes and the richness of the underlying representations.

References

Braje, W. L., Legge, G. E., & Kersten, D. (2000). Invariant recognition of natural objects in the presence of shadows. *Perception, 29*(4), 383-398. doi:10.1068/p3051

Bruce, V., Doyle, T., Dench, N., & Burton, M. (1991). Remembering facial configurations. *Cognition, 38*(2), 109-144. doi:10.1016/0010-0277(91)90049-a

Bruce, V., & Langton, S. (1994). The use of pigmentation and shading information in recognising the sex and identities of faces. *Perception, 23*(7), 803-822. doi:10.1068/p230803

Bruce, V., & Young, A. W. (2012). *Face Perception*. Hove: Psychology Press.

Costen, N. P., Parker, D. M., & Craw, I. (1996). Effects of high-pass and low-pass spatial filtering on face identification. *Percept Psychophys, 58*(4), 602-612. doi:10.3758/bf03213093

Cumming, G. (2012). *Understanding the New Statistics*. New York: Routledge.

DeGutis, J., Wilmer, J., Mercado, R. J., & Cohan, S. (2013). Using regression to measure holistic face processing reveals a strong link with face recognition ability. *Cognition, 126*(1), 87-100. doi:10.1016/j.cognition.2012.09.004

Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G*Power 3: a flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behav Res Methods, 39*(2), 175-191. doi:10.3758/bf03193146

Fisher, K., Towler, J., & Eimer, M. (2016). Effects of contrast inversion on face perception depend on gaze location: Evidence from the N170 component. *Cogn Neurosci, 7*(1-4), 128-137. doi:10.1080/17588928.2015.1053441

Galper, R. E. (1970). Recognition of faces in photographic negative. *Psychonomic Science, 19*(4), 207-208.

Gilad, S., Meng, M., & Sinha, P. (2009). Role of ordinal contrast relationships in face encoding. *Proc Natl Acad Sci U S A, 106*(13), 5353-5358. doi:10.1073/pnas.0812396106

Hayes, A. (1988). Identification of two-tone images; some implications for high- and low-spatial-frequency processes in human vision. *Perception, 17*(4), 429-436. doi:10.1068/p170429

Hayes, T., Morrone, M. C., & Burr, D. C. (1986). Recognition of positive and negative bandpass-filtered images. *Perception, 15*(5), 595-602. doi:10.1068/p150595

Johnston, A., Hill, H., & Carman, N. (1992). Recognising faces: effects of lighting direction, inversion, and brightness reversal. *Perception, 21*(3), 365-375. doi:10.1068/p210365

Kaufmann, J. M., & Schweinberger, S. R. (2012). The faces you remember: caricaturing shape facilitates brain processes reflecting the acquisition of new face representations. *Biol Psychol, 89*(1), 21-33. doi:10.1016/j.biopsycho.2011.08.011

Kemp, R., Pike, G., White, P., & Musselman, A. (1996). Perception and recognition of normal and negative faces: the role of shape from shading and pigmentation cues. *Perception, 25*(1), 37-52. doi:10.1068/p250037

Liu, C. H., Collin, C. A., Burton, A. M., & Chaudhuri, A. (1999). Lighting direction affects recognition of untextured faces in photographic positive and negative. *Vision Res, 39*(24), 4003-4009. doi:10.1016/s0042-6989(99)00109-1

Liu, C. H., Collin, C. A., & Chaudhuri, A. (2000). Does face recognition rely on encoding of 3-D surface? Examining the role of shape-from-shading and shape-from-stereo. *Perception, 29*(6), 729-743. doi:10.1068/p3065

Loftus, G. R., & Harley, E. M. (2005). Why is it easier to identify someone close than far away? *Psychon Bull Rev, 12*(1), 43-65. doi:10.3758/bf03196348

McKelvie, S. J. (1976). The role of eyes and mouth in the memory of a face. *American Journal of Psychology*, 311-323.

Mohr, S., Wang, A., & Engell, A. D. (2018). Early identity recognition of familiar faces is not dependent on holistic processing. *Soc Cogn Affect Neurosci, 13*(10), 1019-1027. doi:10.1093/scan/nsy079

Nasanen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Res, 39*(23), 3824-3833. doi:10.1016/s0042-6989(99)00096-6

O'Toole, A. J., Vetter, T., & Blanz, V. (1999). Three-dimensional shape and two-dimensional surface reflectance contributions to face recognition: an application of three-dimensional morphing. *Vision Res, 39*(18), 3145-3155. doi:10.1016/s0042-6989(99)00034-6

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*(1), 118-127.

Ruiz-Soler, M., & Beltran, F. S. (2006). Face perception: an integrative review of the role of spatial frequencies. *Psychol Res, 70*(4), 273-292. doi:10.1007/s00426-005-0215-z

Russell, R., & Sinha, P. (2007). Real-world face recognition: the importance of surface reflectance properties. *Perception, 36*(9), 1368-1374. doi:10.1068/p5779

Russell, R., Sinha, P., Biederman, I., & Nederhouser, M. (2006). Is pigmentation important for face recognition? Evidence from contrast negation. *Perception, 35*(6), 749-759. doi:10.1068/p5490

Sadr, J., Jarudi, I., & Sinha, P. (2003). The role of eyebrows in face recognition. *Perception, 32*(3), 285-293. doi:10.1068/p5027

Sandford, A., Sarker, T., & Bernier, T. (2018). Effects of geometric distortions, Gaussian blur, and contrast negation on recognition of familiar faces. *Visual Cognition, 26*(3), 207-222.

Sormaz, M., Andrews, T. J., & Young, A. W. (2013). Contrast negation and the importance of the eye region for holistic representations of facial identity. *J Exp Psychol Hum Percept Perform, 39*(6), 1667-1677. doi:10.1037/a0032449

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *Q J Exp Psychol A, 46*(2), 225-245. doi:10.1080/14640749308401045

Vuilleumier, P., Armony, J. L., Driver, J., & Dolan, R. J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nat Neurosci, 6*(6), 624-631. doi:10.1038/nn1057

Wiese, H., Chan, C. Y. X., & Tüttenberg, S. C. (2019). Properties of familiar face representations: Only contrast positive faces contain all information necessary for efficient recognition. *J Exp Psychol Learn Mem Cogn, 45*(9), 1583-1598. doi:10.1037/xlm0000665

Wiese, H., Tüttenberg, S. C., Ingram, B. T., Chan, C. Y. X., Gurbuz, Z., Burton, A. M., & Young, A. W. (2019). A Robust Neural Index of High Face Familiarity. *Psychological Science, 30*(2), 261-272. doi:10.1177/0956797618813572

Young, A. W., & Burton, A. M. (2017). Recognizing Faces. *Current Directions in Psychological Science, 26*(3), 212-217. doi:10.1177/0963721416688114

Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configurational information in face perception. *Perception, 16*(6), 747-759. doi:10.1068/p160747

Appendix

Celebrity Images in Figures 1-4.

Joe Biden 01

Source: https://www.flickr.com/photos/58993040@N07/13978713242

Attribution: U.S. Embassy, Kyiv, Ukraine

Source: https://www.flickr.com/photos/100836534@N04/37487294841

Attribution: LBJ Library, Jay Godwin