Configural face encoding and spatial frequency information

ISABELLE BOUTET University of Montreal, Quebec, Canada

CHARLES COLLIN National Research Council, Ottawa, Ontario, Canada

and

JOCELYN FAUBERT

University of Montreal, Quebec, Canada

Configural relations and a critical band of spatial frequencies (SFs) in the middle range are particularly important for face recognition. We report the results of four experiments in which the relationship between these two types of information was examined. In Experiments 1, 2A, and 2B, the *face inversion effect* (FIE) was used to probe configural face encoding. Recognition of upright and inverted faces and nonface objects was measured in four conditions: a no-filter condition and three SF conditions (low, medium, and high frequency). We found significant FIEs of comparable magnitudes for all frequency conditions. In Experiment 3, discrimination of faces on the basis of either configural or featural modifications was measured under the same four conditions. Although the ability to discriminate configural modifications. We conclude that the band of SF that is critical for face recognition does not contribute preferentially to configural encoding.

Two types of information have been shown to be more important for discrimination and identification of faces than for other complex objects: configural relations and a critical band of spatial frequencies (SFs). With respect to configural relations, several different experimental manipulations have been used to suggest that face recognition depends more heavily on processing of configural or holistic information than does recognition of other objects (for a review, see Maurer, Le Grand, & Mondloch, 2002). Although the words configural and holistic are often used interchangeably in the literature, configural usually refers to the spatial relations between the features of a face (Carey & Diamond, 1994) and holistic to a tendency to process all of the information present in a face (Tanaka & Farah, 1993). Configural information may be more important for faces than for objects because faces are highly homogeneous and because they are recognized at the individual level (e.g., "John's face"). In contrast, recognition of other objects usually takes place at the categorical level (e.g., "a chair" vs. "a house"), at which the presence of one salient feature may be sufficient for discrimination. The "special" status

often attributed to faces may in fact be due to expertise, since observers who are trained in recognizing nonface objects at the individual level appear to use the same type of configural analysis as that usually applied to faces (Diamond & Carey, 1986; Gauthier & Tarr, 1997; Gauthier, Williams, Tarr, & Tanaka, 1998).

Research also suggests that recognition of faces is preferentially supported by a limited range of SFs. Although different experimental paradigms have led to inconsistent results, there is now general agreement that this critical band lies between 6 and 12 cycles per face-width (Bachmann, 1991; Costen, Parker, & Craw, 1996; Fiorentini, Maffei, & Sandini, 1983; Näsänen, 1999). This band of SF is said to be critical because face images that have been filtered and that contain only information within this band will be better recognized than filtered faces containing SFs either below or above this band. In contrast, object recognition is generally unaffected by spatial frequency filtering, even if the task implies recognition at the individual level (Biederman & Kolacsai, 1997).

The reasons for which such a critical band exists remain undetermined, and a number of possibilities are plausible. The advantage of medium frequencies may reflect the selective use of SF information within that range by a specialized face recognition system. Other such specialized systems, such as that devoted to the recognition of letters, may make use of SF information in other ranges (Solomon & Pelli, 1994).

This research was supported by a fellowship from the Canadian Institutes of Health Research (CIHR) to I.B. and CIHR Grant R0010026 to J.F. The authors thank Elizabeth Foley for her translation services. Correspondence concerning this article should be addressed to I. Boutet, Laboratoire de psychophysique et de perception visuelle, Université de Montréal, 3744 Jean-Brilliant, Montreal, Quebec, H3C 3J7 Canada (email: i.boutet@umontreal.ca).

Another related possibility is that the medium range is critical for encoding the configural properties that are essential for face recognition. Wenger and Townsend (2000) examined the relationship between configural encoding and SF information using two types of faces: low-pass and high-pass faces that contained information either above or below a 4.2-cycles/degree cutoff. Participants were asked to discriminate a target face from a distractor when the two differed with respect to their configurations, their parts, or both. Their results indicate that configural changes can be detected from both low-pass and high-pass faces, suggesting that SF information has little influence on configural encoding. However, that study did not fully answer the question of whether the critical band of SF is advantageous for configural encoding, since this band was not examined in isolation.

Nagayama, Yoshida, and Toshima (1995) examined the relationship between configural encoding and SF information using the face inversion effect (FIE). The FIE refers to the finding that face recognition is more significantly impaired by inversion than is recognition of other complex objects (for a review, see Valentine, 1988). There is now considerable evidence that inverted faces are difficult to recognize because inversion disrupts the configural encoding processes applied to upright faces. Behavioral studies that have employed manipulations that vary either the configural or the featural content of faces have shown that inversion has a greater detrimental effect on encoding configural information than on encoding featural information (Freire, Lee, & Symons, 2000; Leder & Bruce, 2000). In addition, many of the effects that illustrate holistic encoding are eliminated or reduced by inversion (Tanaka & Farah, 1993; Tanaka & Sengco, 1987; Thompson, 1980; Young, Hellawell, & Hay, 1987). For example, it has been shown that greater difficulty in recognizing composite faces created by aligning the top half of one familiar face with the bottom half of another familiar face is present only for upright faces. Results from neuropsychological studies also support the notion that the FIE reflects a disruption in the configural encoding of upright faces. Whereas prosopagnosic patients, who are unable to recognize familiar faces as a result of damage to a putative holistic face recognition system, are better at recognizing inverted faces than normals are, Patient C.K., who has object agnosia without prosopagnosia, has more difficulties than normals do in recognizing inverted faces (Moscovitch & Moscovitch, 2000). Together, these results indicate that the FIE reflects encoding of configural information and that inversion has little or no effect on recognition of images that trigger part-based encoding.

In the study by Nagayama et al. (1995), participants had to determine the familiarity of upright and inverted faces in four different frequency conditions (center frequencies of 6.2, 12.4, 24.8, and 49.7 cycles per face-width [c/fw]). Their results showed that all frequency conditions gave rise to a significant FIE. The only observed difference was that this effect was most pronounced for familiar faces at the 6.2-c/fw condition, suggesting that information in the lowest range of the critical band is essential for encoding of configural information. However, this study was limited by the finding that faces containing SFs within the critical band did not produce the expected advantage in recognition performance. In addition, the use of a small number of faces (eight familiar and eight unfamiliar) and the absence of a control condition with a nonface category makes it difficult to evaluate the significance of the reported FIE.

The goal of the experiments reported in this paper was to further investigate the relationship between SF information and configural encoding. In all of the experiments, four frequency conditions were employed: a no-filter condition, a low-frequency condition which 1.25–5 cycles per object width (c/ow) were retained, a medium-frequency condition in which 5–20 c/ow were retained, and a highfrequency condition in which 20–80 c/ow were retained. Both the no-filter and the medium-frequency conditions contained the critical SF band identified in previous studies (Bachmann, 1991; Costen et al., 1996; Fiorentini et al., 1983; Näsänen, 1999).

In Experiments 1 and 2, configural face encoding was examined using the FIE. Previous studies have shown that different SF information may be important for face processing depending on whether the task at hand relies on recognition or discrimination processes (Costen et al., 1996; Hayes, Morrone, & Burr, 1986; Wenger & Townsend, 2000). To address this issue, we employed two different paradigms. An old/new recognition paradigm was used in Experiment 1, and a sequential-matching paradigm was used in Experiment 2. A finding that medium-frequency faces produce the greatest FIE would support the hypothesis that the middle band of SFs is critical for face recognition because this band is essential for configural encoding. Recognition of two nonface categories was also measured to verify that the effects of our different manipulations were specific to faces and not to a difficulty in judging inverted or filtered images.

In Experiment 3, participants discriminated faces that differed with respect to configural or featural information. The same four frequency conditions used in the other experiments were employed, and the faces were shown either upright or inverted. Previous findings have shown that discrimination of configural modifications is more greatly affected by inversion than is discrimination of featural modifications (Freire et al., 2000; Leder & Bruce, 2000). A finding that the medium-frequency condition produced the greatest difference between upright and inverted faces for discrimination of featural modifications but not for discrimination of featural modifications would support the notion that the middle band of SFs is critical for face recognition because this band is essential for configural encoding.

EXPERIMENT 1

Method

Participants. Thirty (13 male) students recruited at Université de Montréal participated in this study. Their ages ranged from 19 to 36 years. All the participants had normal or corrected-to-normal vision.

Apparatus. The participants were tested individually using a Macintosh G3/266 computer with a 21-in. Macintosh color monitor with a refresh screen rate of 75 Hz. The screen was calibrated to linearized luminance values using a Minolta photometer.

Stimuli. Two sets of stimuli were used, one for the *chair* block and one for the *house* block. The chair set consisted of images of 80 faces and 80 chairs. The face images were digitized photographs of males obtained from a database at the University of Essex (http://hpl/essex.ac.uk/projects/vision/allfaces/). The chair images were taken from various websites and from scanned photographs. The house set consisted of images of 120 faces and 120 houses. Half of the face images were of females and half were of males. They were obtained from the Max Planck Institute for Biological Cybernetics in Tübingen, Germany. The house images were obtained from various websites. All the images were converted to a 256 level gray-scale format and cropped to fit in a 13 \times 14 cm window (subtending 13 \times 14 deg of arc at a viewing distance of 57 cm).

Three filtered versions of each stimulus were created as follows. The original images were Fourier transformed into the frequency domain and then multiplied with a pair of 2-D Butterworth filters one high pass and one low pass—of Exponent 5 before being inverse Fourier transformed back into the spatial domain. The frequencies retained by the images were 1.25-5 c/ow for the low-frequency images, 5-20 c/ow for the medium-frequency images, and 20-80 c/ow for the high-frequency images. The low-pass filters were defined by

$$\frac{1}{1+(r/c)^{5}}$$

and the high-pass filters by

$$\frac{1}{1+(c/r)^{5}},$$

where r is the component radius and c the cutoff radius. Prior to filtering, all the images were equated for luminance and root mean square contrast (see the examples in Figure 1 and the image spectra in Figure 2).

Procedure. An old/new recognition paradigm was used. Two blocks were tested: a chair block and a house block. Each block consisted of eight learning/testing segments. Ten images were shown for 6 sec each during each learning segment, with an interstimulus interval of 0.5 sec. In the following testing segment, the 10 learned images were shown randomly intermixed with 10 new ones. On presentation of each image, the participants had to determine whether the image was old or new using the appropriate key. The participants were instructed to take the time needed to give their answers. A break of 1 min was provided after each segment. The images were presented upright in half of the segments and inverted in the other half. Image orientation was the same during learning and testing. Faces were presented in half of the segments and objects in the other half. Presentation order of the segments was randomized. Of the 40 learned images presented across sequences, 10 were not filtered, 10 were in the lowfrequency condition, 10 were in the medium-frequency condition, and 10 were in the high-frequency condition. During testing, the frequency condition of the target was matched with the new distractor images.

For the chair block, 40 faces and 40 chairs were randomly chosen from the corresponding set and used as learning targets. Because many different types of chairs were used (e.g., kitchen chair, office chair), each chair target was matched with a distractor chair with similar physical characteristics. For the house block, 40 faces and 40 houses were randomly chosen from the corresponding set and used as learned targets. Half of the target faces were of males and half were of females. Presentation order of the chair and house blocks was counterbalanced across participants.

Results

Average d' and average reaction times (RTs) for correct responses are illustrated in Figure 3. d' was used as the dependent variable in Experiment 1 because we employed an old/new recognition paradigm that is particularly vulnerable to changes in criterion. This is not the case with the two-alternative forced-choice task used in the other experiments. d' provides a sensitivity measure that takes into account response biases and changes in the internal criterion of the observer. d' was calculated by subtracting an approximated z score for hits from an approximated z score for false alarms using a program designed by Macmillan and Creelman (1991). For each participant, outlier RTs were discarded by taking the mean obtained by each participant in a given condition and discarding RTs that were more than two SDs below or above that mean.

d'. A 2 × 2 × 2 × 4 repeated measures analysis of variance (ANOVA) with block (chair vs. house), image type (face vs. object), orientation (upright vs. inverted), and frequency (no filter vs. low vs. medium vs. high frequency) was performed on the average d' obtained across participants. Neither the effect of block nor any of the interactions involving that factor was significant. Therefore, a 2 × 2 × 4 repeated measures ANOVA was performed on data collapsed across blocks.

The main effects of image type [F(1,29) = 19.90, p < .01], orientation [F(1,29) = 17.49, p < .01], and frequency [F(3,87) = 7.17, p < .01] were significant. The image type × orientation interaction was significant [F(1,29) = 8.03, p = .01]. Contrast analyses revealed that face recognition was significantly affected by inversion [F(1,87) = 20.09, p < .01] but object recognition was not [F(1,87) < 1]. The image type × frequency [F(3,87) = 1.70, p = .17] and the orientation × frequency [F(3,87) = 1.07, p = .37] interactions were not significant. The image type × orientation × frequency interaction was significant [F(3,87) = 3.18, p = .03].

Planned contrast analyses were conducted to test the effect of inversion on face recognition for each frequency condition (see Table 1). A FIE was found for all frequency conditions except the high-frequency condition. The significant image type \times orientation \times frequency interaction can be attributed to the finding that for object recognition, none of the frequency conditions was significantly affected by inversion.

Additional planned contrasts were performed to test whether the medium-frequency condition produced the best recognition performance for upright faces. As can be seen in Table 2, the medium-frequency condition produced a significantly better recognition performance than the high-frequency condition.

RT. The same analyses were performed on the RT data. The main effects of image type [F(1,29) = 33.32, p < .01] and frequency [F(3,87) = 117, p < .01] were significant. The image type × orientation interaction was not significant [F(1,29) < 1]. The image type × frequency [F(3,87) = 27.13, p < .01] and the orientation × frequency [F(3,87) = 21.19, p < .01] interactions were significant. The image type × orientation × frequency interaction was significant [F(3,87) = 10.92, p < .01]. No other effect was significant.

Contrast analyses revealed a significant FIE for the nofilter and the low-frequency conditions (see Table 1).



Figure 1. An example of the images used in Experiments 1 and 2. The no-filter, low-, medium-, and high-frequency images are shown from top to bottom.

Whereas the inversion effect predicts longer RTs for inverted faces, our results indicate that for the medium-frequency condition, RTs were actually much shorter in the inverted condition in comparison with the upright condition. As can be seen in Table 2, a comparison of RTs for upright faces revealed that the high-frequency condition produced slower performance than all of the other conditions.

Discussion

Our results indicate that for upright faces, recognition was best in the no-filter condition, followed by the medium-, low-, and high-frequency conditions. However, performance in the medium-frequency condition was not significantly greater than that in the low-frequency condition. Performance in the medium-frequency condition was nonetheless closest to that in the no-filter condition. Hence, our results provide partial support for the notion that filtered faces containing SF information within the 5–20-c/ow range led to the best recognition performance (Bachmann, 1991; Costen et al., 1996; Fiorentini et al., 1983; Näsänen, 1999). In agreement with previous studies (Biederman & Kalocsai, 1998), object recognition did not differ signifi-



Figure 2. Mean magnitude versus frequency plot of the spatial frequency spectra of the spatially filtered stimuli used in Experiments 1 and 2.

cantly across SF conditions, suggesting that object recognition is supported by a wider range of SFs than face recognition.

The finding that the medium-frequency condition did not produce significantly better performance than the lowfrequency condition warrants further consideration. One difference between Experiment 1 and those studies that identified the medium band as critical for face recognition lies in the use of filtered faces in the learning stage of our experiment. One possibility is that the use of filtered faces was responsible for the lack of a difference between the low- and medium-frequency conditions in Experiment 1. This possibility was explored further in Experiments 2A and 2B, in which learned faces were either unfiltered (Experiment 2A) or filtered (Experiment 2B).

We used the inversion effect to examine the relationship between the SF content of faces and configural encoding. Our assumption was that the presence of a significant inversion effect in a given condition would be an indication that configural encoding was applied to upright faces in that condition. As we expected, our results replicated the FIE, with face recognition being more affected by inversion than object recognition overall. With respect to the different conditions, only the high-frequency condition did not produce a significant FIE. More importantly, the magnitude of the FIE did not differ across the low- and medium-frequency conditions, suggesting that medium SF information does not provide an advantage for configural encoding.

Our RT results were not entirely consistent with the recognition results. Only the low-frequency condition yielded a significant FIE with that measure. Interestingly, upright and inverted objects as well as inverted faces yielded the fastest performance in the medium-frequency condition. Hence, a reversed FIE effect was observed in the mediumfrequency condition for RTs. These results should, however, be viewed with caution, given the low-accuracy performance obtained in all the conditions (ranging between 57% and 69% correct responses). RT is a more sensitive measure when the task is relatively easy and the level of



Figure 3. Mean d' and reaction times (RTs) for correct responses obtained in Experiment 1 for upright and inverted faces and objects in the no-filter, low-frequency (LF), medium-frequency (MF), and high-frequency (HF) conditions. Error bars represent ± 1 SE.

accuracy is high. In difficult tasks, however, accuracy data are more relevant because differences in RT are more easily obscured by variability.

EXPERIMENT 2

Our goal in Experiment 2 was twofold. First, we wished to replicate the results obtained in Experiment 1 using a sequential matching paradigm that would trigger a different response strategy and yield better recognition performance. In Experiment 1, we used an old/new recognition paradigm that relies on memory. In Experiment 2, we used a sequential matching paradigm that relies more heavily on discrimination with little or no role for memory processes. Second, we wished to examine whether the use of filtered learned images was responsible for the lack of a clear mediumfrequency advantage in Experiment 1.

In Experiment 2, a sequential matching procedure was used whereby a target image was quickly presented in the center of the monitor, followed by two test images: the

for Upright and Inverted Faces in Experiment 1			
Frequency Condition	Upright versus Inverted		
d'			
No filter	F = 17.48, p < .01		
Low frequency	F = 7.92, p = .01		
Medium frequency	F = 8.92, p < .01		
High frequency	F < 1		
RT			
No filter	F = 11.42, p < .01		
Low frequency	F = 7.06, p = .01		
Medium frequency	F = 58.81, p < .01		
High frequency	F < 1		

 Table 1

 Contrast Analyses Comparing d' and Reaction Times (RTs)

 for Upright and Inverted Faces in Experiment 1

Note—Degrees of freedom for reported F values are 1 and 87.

sample along with one distractor. The same image manipulations as those used in Experiment 1 were employed, with faces and nonface images presented upright or inverted in four possible frequency conditions. In Experiment 2A, the four conditions were applied only to the test images, and the target was always unfiltered. Experiment 2B was similar to Experiment 1 in that the target was filtered during both learning and testing.

Method

Participants. Twenty-nine (15 male) students participated in Experiment 2A. Their ages ranged from 18 to 46 years. Twenty-eight (8 male) students participated in Experiment 2B. Their ages ranged from 18 to 28 years. All the participants had normal or corrected-to-normal vision. None had participated in Experiment 1.

Apparatus and Stimuli. The apparatus and stimuli were the same as those used in Experiment 1.

Procedure. A sequential matching procedure was used in Experiment 2. One trial consisted of the presentation of one target image at the center of the monitor during learning, followed by a 2-sec mask and then by the presentation of two test images. The test images were presented on the left and right of fixation. The participants selected the test image that matched the learned image as quickly as possible, using the appropriate key. Each session consisted of 80 trials. Faces were presented in half of these trials, and objects were presented in the other half. In Experiment 2A, learned images were either not filtered or filtered in the low-, medium, or high-frequency condition. The test images were shown in the same conditions as the targets. In Experiment 2B, all learned images were intact (no filter), and the test images were either not filtered or filtered in the low-, medium-, or high-frequency condition. Ten trials were run for each image category and for each condition. Presentation order of the image category and conditions was randomized. Recognition accuracy (% correct responses) and RTs were recorded.

In Experiment 2A, the target image was presented for 150 msec at learning. In a pilot for Experiment 2B, we found that recognition performance was at ceiling when the same presentation time was used, suggesting that the matching task was facilitated by the use of filtered images at learning. We therefore reduced the presentation time of the target to 75 msec for Experiment 2B.

Results: Experiment 2A

Average accuracy (percentage of correct responses) and average RTs for correct responses are illustrated in Figure 4. Outlier RTs were discarded in the same fashion as in Experiment 1.

Accuracy. A $2 \times 2 \times 2 \times 4$ repeated measures ANOVA with block (chair vs. house), image type (face vs. object), orientation (upright vs. inverted), and frequency (no filter vs. low vs. medium vs. high frequency) was performed on the average percentage of correct responses obtained across participants. Neither the effect of block nor any of the interactions involving that factor was significant. Therefore, a $2 \times 2 \times 4$ repeated measures ANOVA was performed on data collapsed across blocks.

The main effects of image type [F(1,28) = 80.57, p < .01], orientation [F(1,28) = 24.29, p < .01], and frequency [F(3,84) = 21.76, p < .01] were significant. The image type × orientation interaction was significant [F(1,28) = 5.63, p = .01]. Contrast analyses revealed that face recognition was significantly affected by inversion [F(1,84) = 27.26, p < .01], with upright faces being better recognized than inverted faces (79% and 70%, respectively). As was predicted by inversion [F(1,84) = 3.48, p = .07], with upright and inverted objects yielding 89% and 86% correct responses, respectively. The image type × frequency interaction was significant [F(3,84) = 6.23, p < .01]. No other interaction was significant.

Planned contrast analyses were conducted to test the FIE for each frequency condition (Table 3). All frequency conditions were significantly affected by inversion, and the magnitude of this effect did not differ across conditions.

Table 2
Contrast Analyses Comparing d' and Reaction Times (RTs) Obtained in
Experiment 1 for the Different Frequency Conditions for Upright Faces Only

		- J	0
Frequency Condition	Low Frequency	Medium Frequency	High Frequency
	d	,	
No filter	F = 10.82, p < .01	F = 3.12, p = .08	F = 26.18, p < .01
Low frequency	_	F = 2.32, p = .13	F = 3.33, p = .08
Medium frequency	-	_	F = 11.22, p < .01
	R	Г	
No filter	F < 1	F = 1.58, p = .21	F = 17.14, p < .01
Low frequency	-	F < 1	F = 14.47, p < .01
Medium frequency	-	-	F = 8.29, p = .01

Note—Degrees of freedom for reported F values are 1 and 87.



Figure 4. Mean recognition accuracy (29 participants) and reaction times (RTs; 27 participants) for correct responses obtained in Experiment 2A for upright and inverted faces and objects in the no-filter, low-frequency (LF), medium-frequency (MF), and high-frequency (HF) conditions. Error bars represent ± 1 SE.

Planned contrast analyses were also performed to test whether the medium-frequency condition produced the best recognition performance for upright faces (see Table 4). Of particular relevance for this discussion is the finding that the medium-frequency condition produced the best recognition performance of all the frequency conditions.

RT. Data from 3 participants were discarded because they obtained 0% correct responses in one of the conditions tested. The same analysis as those performed on the accuracy data was performed on the remaining RT data.

The main effects of image type [F(1,25) = 120.63, p < .01] and frequency [F(3,75) = 110.65, p < .01] were significant. The image type × frequency [F(3,75) = 11.13, p < .01] and the orientation × frequency [F(3,75) = 27.85, p < .01]

p < .01] interactions were significant. The image type \times orientation \times frequency interaction was also significant [F(3,75) = 18.31, p < .01]. No other effect was significant. Planned contrast analyses revealed that face recognition was almost significantly affected by inversion [F(1,75) = 3.28, p = .08]. Object recognition was not affected by inversion [F(1,78) < 1].

Planned contrast analyses were conducted to test the FIE for each frequency condition (Table 3). Only the nofilter condition was significantly affected by inversion. As in Experiment 1, a reversed inversion effect was observed in the medium-frequency condition, with RTs to inverted faces being shorter than those to upright faces. Planned contrast analyses were also performed to test whether the

Contrast Analyses Comparing d'and Reaction Times (RTs) for Upright and Inverted Faces in Experiment 2A			
Frequency Condition	Upright versus Inverted		
d	,		
No filter	F = 7.42, p < .01		
Low frequency	F = 13.47, p < .01		
Medium frequency	F = 6.18, p = .02		
High frequency	F = 7.42, p = .01		
R	Г		
No filter	F = 31.42, p < .01		
Low frequency	F < 1		
Medium frequency	F = 134.87, p < .01		

Table 3

Note—Degrees of freedom for reported F values are 1 and 84 for d' and 1 and 75 for RT.

medium-frequency condition produced the best performance for upright faces (see Table 4). In accordance with the recognition data, the medium-frequency condition produced the fastest recognition performance of all the frequency conditions.

Results: Experiment 2B

Average accuracy (percentage of correct responses) and average RTs for correct responses are illustrated in Figure 5. Outlier RTs were discarded in the same fashion as in Experiment 1.

Accuracy. A $2 \times 2 \times 2 \times 4$ repeated measures ANOVA with block (chair vs. house), image type (face vs. object), orientation (upright vs. inverted), and frequency (no filter vs. low vs. medium vs. high frequency) was performed on the average percentage of correct responses obtained across participants. Neither the effect of block nor any of the interactions involving that factor was significant. Therefore, a $2 \times 2 \times 4$ repeated measures ANOVA was performed on data collapsed across blocks.

The main effects of image type [F(1,27) = 74.28, p <.01], orientation [F(1,27) = 34.41, p < .01], and frequency [F(3,81) = 17.23, p < .01] were significant. The image type \times orientation interaction was significant [F(1,27) =7.81, p = .01]. Contrast analyses revealed that face recognition was significantly affected by inversion [F(1,81) =33.70, p < .01], with upright faces being better recognized than inverted faces (82% and 73%, respectively). As was predicted by the FIE, object recognition was not significantly affected by inversion [F(1,81) = 3.43, p = .08], with upright and inverted objects yielding 90% and 87% correct responses, respectively. The image type \times frequency interaction was significant [F(3,81) = 4.70, p < .01]. No other interaction was significant.

As can be seen in Table 5, all conditions produced an FIE. Also relevant to this discussion are the comparisons showing that the medium-frequency condition differed only from the high-frequency condition and that recognition in this condition was the closest to the no-filter condition (Table 6).

RT. Data from 2 participants were discarded because they obtained 0% correct in one of the conditions tested. The same analyses as those performed on the accuracy data were performed on the remaining RT data.

The main effects of image type [F(1,25) = 71.35, p < 71.35].01] and frequency [F(3,75) = 54.43, p < .01] were significant. The image type \times orientation interaction was significant [F(1,25) = 5.80, p = .02]. Planned contrast analyses revealed that face recognition was significantly affected by inversion [F(1,75) = 5.00, p = .04]. Object recognition was not affected by inversion [F(1,75) = 1.38], p = .25]. The image type $\times [F(3,75) = 13.54, p < .01]$ and the orientation \times [F(3,75) = 11.65, p < .01] interactions were significant. The image type \times orientation \times interaction was also significant [F(3,75) = 8.39, p < .01]. No other effect was significant.

As can be seen in Table 5, only the no-filter condition produced a significant FIE. As in Experiments 1 and 2A, a reversed inversion effect was observed in the mediumfrequency condition, with RTs to inverted faces being shorter than those for upright faces.

Planned contrast analyses were also performed to test whether the medium-frequency condition produced the best recognition performance for upright faces (see Table 6). Of particular relevance for this discussion are the comparisons showing that the medium-frequency condition differed only from the high-frequency condition. Since inspection of the results suggested a medium-frequency advantage for the objects as well, additional analyses were conducted for comparison of the medium-frequency condition with the other conditions for upright objects. These analyses revealed that the medium-frequency condition did in fact produce faster RTs for upright objects [low vs. medium: F(1,75) = 6.39, p = .01; high vs. medium: F(1,75) =7.01, p < .01].

Discussion: Experiments 2A and 2B

The results of Experiments 2A and 2B are consistent with the notion that face recognition relies on a critical SF band in the middle range. In Experiment 2A, the mediumfrequency condition produced significantly greater recognition accuracy than the other two conditions. In Experiment 2B, as in Experiment 1, the medium-frequency condition produced the best recognition performance in the sense that the results in this condition were always the closest to those in the no-filter condition. However, recognition accuracy was not significantly different in the lowversus the medium-frequency condition in Experiments 1 and 2B. The use of filtered versus unfiltered learned faces seems to be the most likely explanation for this discrepancy. Whereas unfiltered faces were used during learning in Experiment 2A, filtered faces were used during learning in Experiments 1 and 2B. Moreover, previous studies that have identified the medium-frequency band as being critical for face recognition used unfiltered faces at learning (Bachmann, 1991; Costen et al., 1996; Fiorentini et al., 1983).

Closer inspection of the data indicates that the use of filtered learned images did not really affect performance in the medium-frequency condition but, rather, improved recognition in the low-frequency condition. This improvement was responsible for the lack of a significant difference between the low- and medium-frequency conditions in Experiments 1 and 2B. There is less information in a



Figure 5. Mean recognition accuracy (28 participants) and reaction times (RTs; 28 participants) for correct responses obtained in Experiment 2B for upright and inverted faces and objects in the no-filter, low-frequency (LF), medium-frequency (MF), and high-frequency (HF) conditions. Error bars represent ± 1 SE.

band of low SFs than in an equivalent band of higher SFs. Therefore, an unfiltered learned face and a low-frequency test face have less information in common than an unfiltered learned face and a medium- or high-frequency test face. This could explain why performance was much better in the low-frequency condition when filtered learned images were used than when unfiltered images were used.

The presentation time used in Experiment 2B was much shorter than that used in Experiment 2A, and yet overall recognition performance was equivalent in these two experiments. The absence of interference from other bands may explain why having a perceptual match between the SF information present in a target at learning and at testing improves performance. This interpretation is consistent with the finding that the degree of SF overlap between learned and test faces is just as important as the band of SF present for accurate face recognition (Liu, Collin, Rainville, & Chaudhuri, 2000).

All frequency conditions were significantly affected by inversion in Experiments 2A and 2B, whereas only the low- and medium-frequency conditions were significantly affected by inversion in Experiment 1. Given the general pattern of results observed in these experiments, we interpret the lack of an FIE for the high-frequency condition as a spurious finding subject to the collection of further evidence. Otherwise, our results indicate that SF filtering has little influence on the FIE. Since the FIE served as an indicator of configural encoding, we conclude that although

Contrast Analyses Comparing Accuracy Performance and Reaction Times (RTs) Obtained in Experiment 2A for the Different Frequency Conditions for Upright Faces Only				
Frequency Condition	Low Frequency	Medium Frequency	High Frequency	
Accuracy				
No filter	F = 25.92, p < .01	F = 5.60, p = .02	F = 33.66, p < .01	
Low frequency	-	F = 7.41, p < .01	F < 1	
Medium frequency	-	-	F = 11.79, p < .01	
RT				
No filter	F = 27.34, p < .01	F = 10.66, p < .01	F = 154.69, p < .01	
Low frequency	_	F = 3.85, p = .05	F = 51.97, p < .01	
Medium frequency	-	_	F = 87.12, p < .01	

Table 4

Note—Degrees of freedom for reported F values are 1 and 87 for accuracy and 1 and 75 for RT.

medium SF information does play a critical role in face recognition, it is not more important than other SF information for configural encoding.

Whereas object recognition did not significantly differ across SF conditions in Experiment 1, there was a tendency for medium-frequency objects to be best recognized in Experiment 2A and a clear advantage for that frequency condition in Experiment 2B. This advantage was also observed in the RT data, given that RTs were shortest for the medium-frequency condition for upright and inverted objects in all the experiments. Short RTs were also observed in the medium-frequency condition for inverted faceshence the finding of a reversed FIE for that condition. Possible explanations for this finding are provided in the General Discussion section.

In conclusion, the results of Experiments 1 and 2 suggest that although the medium-frequency band is critical for face recognition, it is not critical for encoding configural information. In Experiment 3, we further examined this hypothesis using a more precise measure of configural encoding. We employed a task in which faces were discriminated on the basis of either featural or configural information (Freire et al., 2000; Leder, 1996; Leder & Bruce, 2000). The advantage of this task is that the requirements are clear: There is no memory involved, and, unlike the case of the FIE tested in Experiments 1 and 2, there is no uncertainty as to whether configural encoding takes place at learning or at testing. Discrepancies with regard to the use of filtered or unfiltered learned faces are also avoided.

EXPERIMENT 3

In Experiment 3, we further examined the relationship between SF information and configural face encoding using a more direct test of configural encoding. Participants were asked to perform a discrimination task on faces that differed with respect to either configural or featural information. Eight baseline faces were used, and nine modified versions of each baseline face were created either by moving the location of the eyes, the nose, or the mouth or by exchanging one of these features with the corresponding feature of another face. The same four frequency conditions as those used in Experiments 1 and 2 were applied to each baseline face and to its modified versions. The participants were shown two modified faces or one baseline and one modified face, and were asked to determine whether they were the same or different.

Upright and inverted conditions were included for evaluation of the efficiency of our manipulations. On the basis of the findings of Freire et al. (2000) and Leder and Bruce (2000), we predicted that discrimination of configural information would be more impaired by inversion than discrimination of featural information for the unfiltered faces. As for the frequency conditions, a finding that mediumfrequency faces produce the best discrimination performance for configural modifications but not for featural modifications would support the hypothesis that the medium band of SF is essential for configural encoding.

Method

Participants. Twenty-eight (7 male) students participated in this experiment. Their ages ranged from 18 to 28 years. All the participants had normal or corrected-to-n ormal vision. None had participated in any of the previous experiments.

Apparatus. The apparatus was the same as that used in Experiment 1.

Stimuli. Eight baseline faces were created using the Adobe Photoshop 6.0 software in the following fashion. Fifty original gray-

Table 5
Contrast Analyses Comparing d' and Reaction Times (RTs)
for Upright and Inverted Faces in Experiment 2B

F			
Frequency Condition		Upright versus Inverted	
	d'		
No filter		F = 5.80, p = .01	
Low frequency		F = 10.56, p < .01	
Medium frequency		F = 13.04, p < .01	
High frequency		F = 4.19, p = .04	
	RT		
No filter		F = 7.0, p = .01	
Low frequency		F = 1.50, p = .23	
Medium frequency		F = 101.65, p < .01	
High frequency		F < 1	

Note—Degrees of freedom for reported F values are 1 and 81 for d' and 1 and 75 for RT

Table 6
Contrast Analyses Comparing Accuracy Performance
and Reaction Times (RTs) Obtained in Experiment 2B
for the Different Frequency Conditions for Upright Faces Only

Frequency Condition	Low Frequency	Medium Frequency	High Frequency
	Accu	racy	
No filter Low frequency	F = 5.80, p = .02	F = 1.75, p = .19 F = 1.17, p = .28	F = 30.66, p < .01 F = 9.79, p < .01
Medium frequency	– R'	– Г	F = 17.48, p < .01
No filter	F < 1	F = 2.52, p < .12	F = 40.57, p < .01
Low frequency	-	F = 1.04, p = 31	F = 33.65, p < .01
Medium frequency	-	-	F = 22.85, p < .01

Note—Degrees of freedom for reported F values are 1 and 87 for accuracy and 1 and 75 for RT.

scale face photographs taken from those used in the previous experiments were cropped so that the eyes, noses, and mouths could be removed from their external features. Using this set of internal and external features, eight different baseline faces were created using different external feature sets and internal features from the original set. Features used to create a given baseline face were always from different original faces to avoid the creation of more realistic-looking faces. The size of the external features was averaged across the baseline faces such that all baseline faces were 7.5 cm in width and 11.5 cm in height (subtending 7.5×11.5 deg of arc at a viewing distance of 57 cm). Nine versions of the eight baseline faces were then created: six for the configural modifications and three for the featural modifications (see Figure 6). Featural modifications were made by swapping the eyes, nose, and mouth of the baseline face with that of a new face. Again, the resulting faces contained sets of features that all came from different original faces. Configural modifications were made by moving the eyes, nose, or mouth up or down by 5 mm.

All baseline and modified faces were filtered using the same procedure as in the other experiments, producing no-filter, low-, medium-, and high-frequency conditions for each baseline face and its modified versions.

Procedure. A simultaneous discrimination paradigm was used. For each trial, two faces were shown side by side with the right face being higher than the other to ensure that the participants could not use an imaginary horizontal line to detect configural modifications. The two faces were the same in half of the trials and different in the other half. Extensive pilot work with varied amounts of configural and featural modifications indicated that unless extensive changes that rendered the faces grotesque and nonrealistic had been made, the participants' performance was near chance in a speeded discrimination task. We therefore decided to have the participants give their answers after a fixed presentation time of 3.5 sec. The faces were therefore shown for 3.5 sec, after which a window filled with noise appeared during which participants gave their answers by pressing the appropriate key. Negative feedback for wrong answers was provided using a tone.

Four baseline faces were randomly chosen for the experiment. For each baseline face, 24 trials were used for each frequency condition. The faces were upright for 12 of these trials and inverted for the other 12 trials. Of the 12 trials for each orientation, 6 were same trials and 6 were different trials. The 6 same trials consisted of the presentation of six modified versions of the baseline face: three versions with featural changes and three versions with configural changes (with features moved either up or down). The 6 different trials consisted of the presentation of the baseline face along with one of the six modified versions used in the same trials. A total of 384 trials were shown in the experiment: 4 faces \times 24 trials per face \times 4 frequency conditions. The order of testing was randomized across conditions.

Testing began with a set of instructions in which the participants were shown one example of each baseline face with six modified versions (one for each configural change with the critical feature being moved either up or down, and one for each featural change). The types of modifications made were explained to the participants, who then familiarized themselves with these modifications by scanning through the baseline face and its six modified versions. This process was repeated with the eight baseline faces. This was followed by a practice of 24 trials. In the practice, an unfiltered baseline face was used. This face was not presented in the main experiment. Presentation time was unlimited for the practice, and the participants had to obtain a performance of at least 85% correct before moving on to the experiment. In the experiment, three breaks were given during the presentation of the 384 trials.

Results

A 2 × 4 × 2 repeated measures ANOVA with orientation (upright vs. inverted), frequency (no filter vs. low vs. medium vs. high frequency), and modification type (configural vs. featural) was performed on percentage of correct responses for the discrimination task (Figure 7). The main effects of orientation [F(1,27) = 49.39, p < .01], frequency [F(3,81) = 61.57, p < .01], and modification type [F(1,27) = 10.81, p < .01] were significant. The orientation × frequency [F(3,81) = 3.26, p = .03], orientation × modification type [F(1,27) = 16.76, p < .01], and frequency × modification type [F(3,81) = 15.18, p < .01] interactions were significant. As was expected, inversion impaired discrimination of configural changes to a greater extent than discrimination of featural changes. No other effect was significant.

Planned contrast analyses were performed to test the predictions outlined in the introduction. Discrimination of configural and featural modifications was examined across all frequency conditions (see Table 7). The results indicate that the medium-frequency condition produced the best discrimination performance for both configural and featural modifications. Hence, there is no evidence that medium SF information provides an advantage for configural encoding.



Figure 6. An example of a baseline face (top row) and the configural (middle row) and featural (bottom row) differences created. The configural differences illustrated in the middle row represent only one possibility for each feature, with the left face illustrating the eyes moved up, the center face illustrating the nose moved down, and the right face illustrating the nose moved up. Three other configural differences not illustrated in the figure were also used in the experiment: eyes moved down, nose moved up, and mouth moved down. Only the no-filter condition is shown here.

Discussion

The goal of Experiment 3 was to provide a more direct assessment of the relationship between SF information and configural encoding processes. A discrimination task was used whereby the participants had to determine whether faces that differed with respect to either configural information or featural information were the same or different. First, the results supported the notion that encoding of configural information is more greatly impaired by inversion than encoding of featural information. Second, the medium-frequency condition produced better performance than all the other frequency conditions for both configural and featural modifications for upright faces. These results are consistent with the notion that SF in the middle range is critical for face processing. However, we failed to find an advantage for detection of configural versus featural changes for medium-frequency upright faces, which suggests that medium SF information is not essential for configural encoding. This finding is consistent with the results of Experiments 1 and 2, in which the mediumfrequency condition did not produce a greater FIE than did the other frequency conditions.

GENERAL DISCUSSION

In the experiments reported in this paper, we examined the relationship between SF information and configural encoding. A band of SFs in the middle range that is critical for face recognition had been identified in previous studies. However, the reasons for which this band of SFs produces better recognition than other SF information remain to be determined. One possibility is that the medium band of SF is critical for face recognition because this band best conveys configural information. Four experiments were conducted to investigate this hypothesis. In Experiments 1 and 2, the FIE was used to probe configural encoding. An FIE is observed when recognition of faces is more greatly impaired by inversion than recognition of



Figure 7. Mean discrimination accuracy (28 participants) obtained in Experiment 3 for featural and configural modifications and upright and inverted faces. Error bars represent ± 1 SE. LF, low frequency, MF, medium frequency; HF, high frequency.

nonface objects. Two different recognition paradigms as well as filtered and unfiltered learned faces were used. All images were presented in one of four frequency conditions: a no-filter condition, a low-frequency condition that retained 1.25–5 c/ow, a medium-frequency condition that retained 5–20 c/ow, and a high-frequency condition that retained 20-80 c/ow. Both the no-filter and the medium-frequency conditions contained the critical SF band identified in previous studies.

The results of Experiments 1 and 2 indicate that SF filtering has little influence on configural encoding as measured by the FIE. Nagayama et al. (1995) reported that all the frequency conditions they used (6.2, 12.4, 24.8, and 49.7 c/fw) gave rise to a significant FIE, the effect being most pronounced for the 6.2-c/fw condition. Their study suggested that configural face encoding is supported by a wide range of SFs, with SFs in the lowest range of the critical band being the most important. However, their study was limited by the finding that faces with SF within the critical band did not produce the expected advantage in recognition performance despite the use of nonfiltered faces in the learning stage of the experiment. Moreover, the use of a small number of faces (eight familiar and eight unfamiliar) and the absence of a control condition with a nonface category makes it difficult to evaluate the significance of the reported FIE.

Our finding that all three frequency conditions produced an FIE was consistent with the results of Experi-

Table 7Contrast Analyses Comparing Discrimination of Configuraland Featural Changes for the Different Frequency Conditionsfor Upright Faces Only in Experiment 3			
Frequency Condition	Low Frequency	Medium Frequency	High Frequency
	Configura	l Changes	
No filter	F = 94.35, p < .01	F < 1	F = 12.65, p < .01
Low frequency	-	F = 88.41, p < .01	F = 37.91, p < .01
Medium frequency	-	-	F = 10.54, p < .01
	Featural	Changes	
No filter	F = 187.66, p < .01	F < 1	F = 5.36, p = .02
Low frequency	_	F = 181.82, p < .01	F = 129.60, p < .01
Medium frequency	_	_	F = 4.41, p = .04

Note-Degrees of freedom for reported F values are 1 and 81.

ment 3. In that experiment, the participants were asked to discriminate two simultaneously presented faces on the basis of configural or featural modifications. As was predicted by the results of previous experiments (Freire et al., 2000; Leder & Bruce, 2000), we found that inversion had a greater detrimental effect on discrimination of configural modifications than on discrimination of featural modifications. Moreover, the medium-frequency condition produced the best discrimination performance for both types of modifications. Together, these results further suggest that the medium band of SF is critical for face processing in general and yet does not preferentially trigger configural encoding.

Our results also suggest that the medium band of SF is critical for face recognition when unfiltered learned faces are used but not when filtered learned faces are used. Whereas low- and medium-frequency faces yielded equivalent recognition performance when unfiltered learned faces were used, medium-frequency faces yielded the best recognition performance when filtered learned faces were used. One possible explanation for this finding is based on the notion that there is less information in a band of low SFs than in an equivalent band of higher SFs. As a result, an unfiltered learned face and a low-frequency test face have less information in common than an unfiltered learned face and a medium- or high-frequency test face. This could explain why recognition was much better in the low-frequency condition when filtered learned images were used than when unfiltered images were used.

If this explanation is correct, then one might ask why high-frequency faces produce lower recognition performance irrespective of whether the learned faces are filtered or unfiltered. Low performance with high-frequency faces can be attributed to the relatively low contrast of these faces. Indeed, one could hypothesize that the mediumfrequency advantage arises because medium-frequency test faces contain both a sufficient amount of information in common with unfiltered learned faces and a sufficiently high contrast to produce the best recognition performance. Although this explanation is speculative, it is consistent with the finding that SF overlap between learned and test images is just as important as SF band for face recognition (Liu et al., 2000). It is also consistent with the finding that the medium band yields the highest signal-to-noise thresholds for matching unfiltered faces with filtered faces (Näsänen, 1999).

The hypothesis that a sufficient amount of overlapping information and contrast is partially responsible for the medium frequency advantage for faces implies that this advantage should also apply to nonface objects. Although previous reports have suggested that recognition of nonface objects is equally supported by different spatial frequency information (e.g., Biederman & Kolacsai, 1998), the results of our experiments do not entirely support this notion. As can be seen in Figures 3 and 4, there was a tendency for medium-frequency objects to produce the best recognition performance. Our finding is perhaps more significant, since recognition of faces and that of nonface objects were directly compared and the level of categorization at which this recognition took place was equated.

Finally, inverted faces and upright and inverted objects produced the shortest RTs for the medium-frequency condition in Experiments 1 and 2. Given that short RTs were not associated with poor recognition performance in all three conditions, speed-accuracy trade-offs are unlikely to have been responsible for this result. A possible explanation for this finding is based on the notions that recognition of inverted faces and objects relies on the identification and detection of a few salient features during learning and testing, and that these features are best extracted when medium SFs are present. If this is the case, then one would find a medium-frequency advantage for detection of these features for inverted faces and for upright and inverted objects. When asked to recognize one of these object categories, participants would more quickly detect the presence of the salient features in the mediumfrequency condition than in the other conditions—hence the short RTs that we have observed in our experiments. However, participants may quickly but wrongly choose a distractor in a speeded task if the distractor contains features that are highly similar to those identified in the target at learning.

Hence, homogeneity between targets and distractors may determine whether a fast detection of salient features results in accurate recognition. Considering that the faces we used were more homogeneous than the objects, this could explain why short RTs were not accompanied by good recognition performance for the inverted faces in the medium-frequency condition. This interpretation would also be in line with the finding that upright faces did not produce the short RTs observed in the medium-frequency condition with inverted faces. Indeed, upright faces are believed to be encoded on the basis of configural information rather than featural information.

In conclusion, the results of our experiments provide evidence that the advantage held by encoding of the medium band of SF in face recognition is not related to configural encoding. Rather, these two processes appear to contribute independently to our face recognition abilities, perhaps because they operate at different stages in the process of recognizing a face. Our results also suggest that processing of featural information is fastest for medium SF information. Further experiments are necessary to determine how such a process might operate and how it would be integrated with other low- and high-level mechanisms that are critical for recognition of upright faces.

REFERENCES

- BACHMANN, T. (1991). Identification of spatially quantised tachistoscopic images of faces: How many pixels does it take to carry identity? *European Journal of Cognitive Psychology*, 3, 85-103.
- BIEDERMAN, I., & KALOCSAI, P. (1997). Neurocomputational bases of object and face recognition. *Philosophical Transactions of the Royal Society of London*, 352, 1203-1219.

- CAREY, S., & DIAMOND, R. (1994). Are faces perceived as configurations more by adults than by children? *Visual Cognition*, 1, 253-274.
- COSTEN, N. P., PARKER, D. M., & CRAW, I. (1996). Effects of high-pass and low-pass spatial filtering on face identification. *Perception & Psychophysics*, 58, 602-612.
- DIAMOND, R., & CAREY, S. (1986). Why faces are and are not special: An effect of expertise. *Journal of Experimental Psychology: General*, 115, 107-117.
- FIORENTINI, A., MAFFEI, L., & SANDINI, G. (1983). The role of high spatial frequencies in face perception. *Perception*, 12, 195-201.
- FREIRE, A., LEE, K., & SYMONS, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception*, 29, 159-170.
- GAUTHIER, I., & TARR, M. J. (1997). Becoming a "greeble" expert: Exploring mechanisms for face recognition. *Vision Research*, 37, 1673-1682.
- GAUTHIER, I., WILLIAMS, P., TARR, M. J., & TANAKA, J. (1998). Training "greeble" experts: A framework for studying expert object recognition processes. *Vision Research*, 38, 2401-2428.
- HAYES, T., MORRONE, M. C., & BURR, D. C. (1986). Recognition of positive and negative bandpass-filtered images. *Perception*, 15, 595-602.
- LEDER, H. (1996). Line drawings of faces reduce configural processing. *Perception*, 25, 355-366.
- LEDER, H., & BRUCE, V. (2000). When inverted faces are recognized: The role of configural information in face recognition. *Quarterly Journal of Experimental Psychology*, **53A**, 513-536.
- LIU, C. H., COLLIN, C. A., RAINVILLE, S. J., & CHAUDHURI, A. (2000). The effects of spatial frequency overlap on face recognition. *Journal* of Experimental Psychology: Human Perception & Performance, 26, 956-979.
- MACMILLAN, N. A., & CREELMAN, C. D. (1991). Detection theory: A user's guide. New York: Cambridge University Press.

- MAURER, D., LE GRAND, R., & MONDLOCH, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6, 255-260.
- MOSCOVITCH, M., & MOSCOVITCH, D. A. (2000). Super face-inversion effects for isolated internal or external features, and for fractured faces. *Cognitive Neuropsychology*, **17**, 201-219.
- NAGAYAMA, R., YOSHIDA, H., & TOSHIMA, T. (1995). Interrelationship between the facial expression and familiarity: Analysis using spatial filtering and inverted presentation. *Japanese Journal of Psychology*, 66, 327-335.
- Näsänen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research*, **39**, 3824-3833.
- SOLOMON, J. A., & PELLI, D. G. (1994). The visual filter mediating letter identification. *Nature*, 369, 395-397.
- TANAKA, J. W., & FARAH, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology*, 46A, 225-245.
- TANAKA, J. W., & SENGCO, J. A. (1997). Features and their configuration in face recognition. *Memory & Cognition*, 25, 583-592.
- THOMPSON, P. (1980). Margaret Thatcher—A new illusion. *Perception*, 9, 483-484.
- VALENTINE, T. (1988). Upside-down faces: A review of the effect of inversion upon face recognition. *British Journal of Psychology*, 79, 471-491.
- WENGER, M. J., & TOWNSEND, J. T. (2000). Spatial frequencies in shortterm memory for faces: A test of three frequency-dependent hypotheses. *Memory & Cognition*, 28, 125-142.
- YOUNG, A. W., HELLAWELL, D., & HAY, D. C. (1987). Configural information in face perception. *Perception*, 16, 747-759.

(Manuscript received November 2, 2002; revision accepted for publication March 25, 2003.)