

Faces in the mist: illusory face and letter detection

Cory A. Rieth

Department of Psychology, University of California, San Diego, La Jolla, CA 92093-0109, USA;
e-mail: crieth@ucsd.edu

Kang Lee

Department of Psychology, University of California, San Diego, La Jolla, CA 92093-0109, USA, and
Institute of Child Study, University of Toronto, Toronto, ON M5R 2X2, Canada; e-mail: kang.lee@utoronto.ca

Jiangang Lui

Department of Biomedical Engineering, School of Computer and Information Technology, Beijing Jiaotong
University, Beijing 100044, China; e-mail: liujg@fingerpass.net.cn

Jie Tian

School of Life Science and Technology, Xidian University, Xi'an 710071, China, and Institute of
Automation, Chinese Academy of Sciences, P. O. Box 2728, Beijing 100190, China; e-mail: tian@ieee.org

David E. Huber

Department of Psychology, University of California, San Diego, La Jolla, CA 92093-0109, USA;
e-mail: dhuber@ucsd.edu

Received 20 January 2011, in revised form 26 May 2011; published online 15 August 2011

Abstract. We report three behavioral experiments on the spatial characteristics evoking illusory face and letter detection. False detections made to pure noise images were analyzed using a modified reverse correlation method in which hundreds of observers rated a modest number of noise images (480) during a single session. This method was originally developed for brain imaging research, and has been used in a number of fMRI publications, but this is the first report of the behavioral classification images. In Experiment 1 illusory face detection occurred in response to scattered dark patches throughout the images, with a bias to the left visual field. This occurred despite the use of a fixation cross and expectations that faces would be centered. In contrast, illusory letter detection (Experiment 2) occurred in response to centrally positioned dark patches. Experiment 3 included an oval in all displays to spatially constrain illusory face detection. With the addition of this oval the classification image revealed an eyes/nose/mouth pattern. These results suggest that face detection is triggered by a minimal face-like pattern even when these features are not centered in visual focus.

Keywords: vision, face perception, reverse correlation, letter perception, top down, false detection.

1 Introduction

When looking at randomly positioned lunar craters, arbitrary wisps of cloud, or a scattered pile of rocks, one often has the impression of seeing a face. Illusory detection is termed “pareidolia”, and popular examples include the Man on the Moon, the face in the Cydonia region of Mars, and the faces of numerous religious icons in toasted food (Svoboda 2007). Providing many other examples, a blog called “faces in places” catalogs everyday objects that look like faces (<http://facesinplaces.blogspot.com/>). Although these objects are not mistaken for actual faces, they evoke the percept of a face in a compelling manner. In contrast to false face recognition, which is the mistaking of one face for another, illusory face detection is the reported detection of a face when no face image exists.

The process of falsely recognizing visible faces has been heavily investigated, particularly as it relates to false identifications during police lineups (Lindsay et al 1991), but the processes underlying face detection, especially illusory face detection, are not well understood at the behavioral level (Lewis and Edmonds 2003). Nevertheless, recent research has examined illusory detection and similar paradigms using neuroimaging. The fusiform face area (FFA) is an area in the inferotemporal cortex that exhibits greater responses for faces than most other objects (Kanwisher et al 1997). The FFA is not only preferentially active when viewing faces (e.g., bottom-up processing) but is also active during top-down face processing—

for example, when imagining faces (O'Craven and Kanwisher 2000), when anticipating faces (Esterman and Yantis 2010), and when interpreting bistable images as faces (Andrews et al 2002; Hasson et al 2001). However, these tasks either present a coherent bottom-up signal (e.g., a bistable image that is easily interpreted as a face) or rely on the introspective processes of the observer (e.g., instructions to imagine a face). To address these limitations, we developed a new technique for eliciting illusory face detection that provides completely ambiguous noise, thus providing a relatively pure measure of top-down processing without reliance on introspection. Using this paradigm, we examined the neural correlates of both illusory face detection and illusory letter detection (Li et al 2009, 2010; Liu et al 2010, 2011; Zhang et al 2008). In the current experiments we used this same paradigm to provide the first behavioral account of the visual template that observers likely used while experiencing illusory detection in these fMRI experiments.

In this illusory detection paradigm faces or letters become progressively more difficult to perceive throughout the experiment until, unbeknownst to observers, every trial presents a pure noise image. Using observers' detection responses in combination with measures of neural activity, we found that the FFA is selectively active when detecting illusory faces (Zhang et al 2008). Additionally, we identified a distributed network of brain areas involved in illusory face detection (Li et al 2009, 2010). This network largely overlaps with regions that are active when viewing actual faces (Fairhall and Ishai 2007; Ishai 2008), except that some earlier visual areas are absent, such as the occipital face area. For illusory letter detection a similar distributed network of brain regions was identified (Liu et al 2010, 2011). Missing from these investigations is specification of the spatial attributes of a noise image that tend to promote illusory detection of faces or letters. In other words, these studies tell us which brain regions are involved in illusory detection, but they do not tell us what visual properties activated those brain regions. The current experiments answer this question by replicating the illusory face and letter detection experiments outside the scanner. Using a technique called "reverse correlation" (Dayan and Abbott 2001), we computed correlations between detection responses and the noise images. In this way, we determined the aspects of the noise images that trigger illusory detection for faces and letters.

1.1 Outline of the current experiments

Faces and letters are two very different classes of objects in terms of their visual properties. For instance, any two faces are much more visually similar to each other than are any two letters. In light of these differences a comparison between illusory detection of faces and letters may seem unwise. However, beyond the fact that our paradigm has already been used in fMRI experiments with these two classes of stimuli there are other advantages to this comparison. For instance, faces and letters are two classes of visual stimuli that have been heavily studied, and most adults are expert at identifying both. Furthermore, letters and faces appear to be processed in different parts of the brain; and, correspondingly, there are distinct neuropsychological deficits for each stimulus class (e.g., prosopagnosia and dyslexia). Of critical importance for the current experiments, the brain regions that are specialized for faces and letters appear in opposite cortical hemispheres: the FFA typically resides in the right hemisphere (Kanwisher et al 1997), while the visual word form area typically resides in the left hemisphere (McCandliss et al 2003). This cortical laterality may exhibit itself in the form of spatial laterality differences for the noise patterns that promote illusory detection of faces and letters.

Experiments 1 (illusory face detection) and 2 (illusory letter detection) are direct replications of our fMRI experiments (Li et al 2009, 2010; Liu et al 2010, 2011; Zhang et al 2008) except that the current experiments collected data from many more observers to determine the spatial attributes of the noise images that tend to promote illusory detection. Experiment

3 tested whether additional bottom-up constraints could eliminate the noncentrality of face detection by including an oval in the noise images.

Reverse correlation techniques were originally developed to study the spatial characteristics of neural receptive fields (Dayan and Abbott 2001; Jones and Palmer 1987; Movshon et al 1978; Ringach and Shapley 2004). To determine the spatial characteristics that promote detection responses, a correlation is calculated for each observer at each pixel position for the relationship between luminance (e.g., gray scale value as determined by adding pixel noise to an image) and whether the observer detected a target on the corresponding trial (in the case of a neuron, detection is determined by spike rate). This analysis creates a “classification image” (CI) showing regions of the noise images where luminance correlated with responses. To study illusory detection, our design and analysis deviated from typical behavioral reverse correlation experiments by using all noise images, averaging detection responses over a large number of observers, using complex but random images, and establishing the expectancy of targets through training blocks instead of through inclusion of targets on a subset of trials. Figure 1 provides a graphical summary of the typical techniques used and an overview of the design of the current study. This variant is new, and our results demonstrate its ability to identify the biases and expectations that underlie detection of an object class.

Most reverse correlation experiments based on behavioral responses present a true target image combined with independent pixel noise such that detection is influenced by the particular noise pattern on each trial (e.g., Ahumada 1996; Gold et al 2000; Kontsevich and Tyler 2004). This is also true for studies using neural measurements in humans (Smith et al 2008, 2009). For example, to find the visual features necessary to detect a smile, observers were shown the *Mona Lisa* combined with noise and asked to rate her emotional expression on a scale from sad to happy (Kontsevich and Tyler 2004). Using targets combined with noise to investigate illusory detection of faces and letters is problematic for two reasons. First, if a target is present, detection is not illusory. Second, a detection task for the same target image on every trial may differ from illusory detection for any item from the object class. To study top-down expectations for detection of faces and letters as object classes without contamination from bottom-up information, we needed a task that did not include actual target images during the critical experimental trials. We also needed observers to expect any face rather than a particular face (or letter). Therefore, we used a training phase that presented a variety of different centered face/letter images mixed with noise, with these targets becoming progressively more difficult to detect. Eventually, target images were omitted entirely, and only noise images were used for the remainder of the experiment. This training established top-down expectations about what class of object to expect (faces in Experiments 1 and 3, and letters in Experiment 2) and where to expect an exemplar from that object class (in the center of the image). Although a previous study has used noise-only images to create a behavioral CI (Experiment 1 of Gosselin and Schyns 2003), this is the first time that this has been done with detection of an object class rather than a specific target (e.g., the *Mona Lisa's* smile) or a specific exemplar (e.g., the letter “s”). This is an important distinction. With a known target observers may adopt a search strategy that is unique to that target (i.e., reliance on specific low-level visual features). In contrast, for the detection of an unknown target from an object class, observers are forced to rely on class general top-down expectations.

The presence of visible targets appearing throughout the experiment helps observers stay focused on the detection task across the thousands of trials used in a typical reverse correlation experiment. However, illusory detection for an object class would be contaminated by including visible targets on a subset of trials because the preferred image would be heavily biased towards the specific form of the most recently seen target. Therefore, to

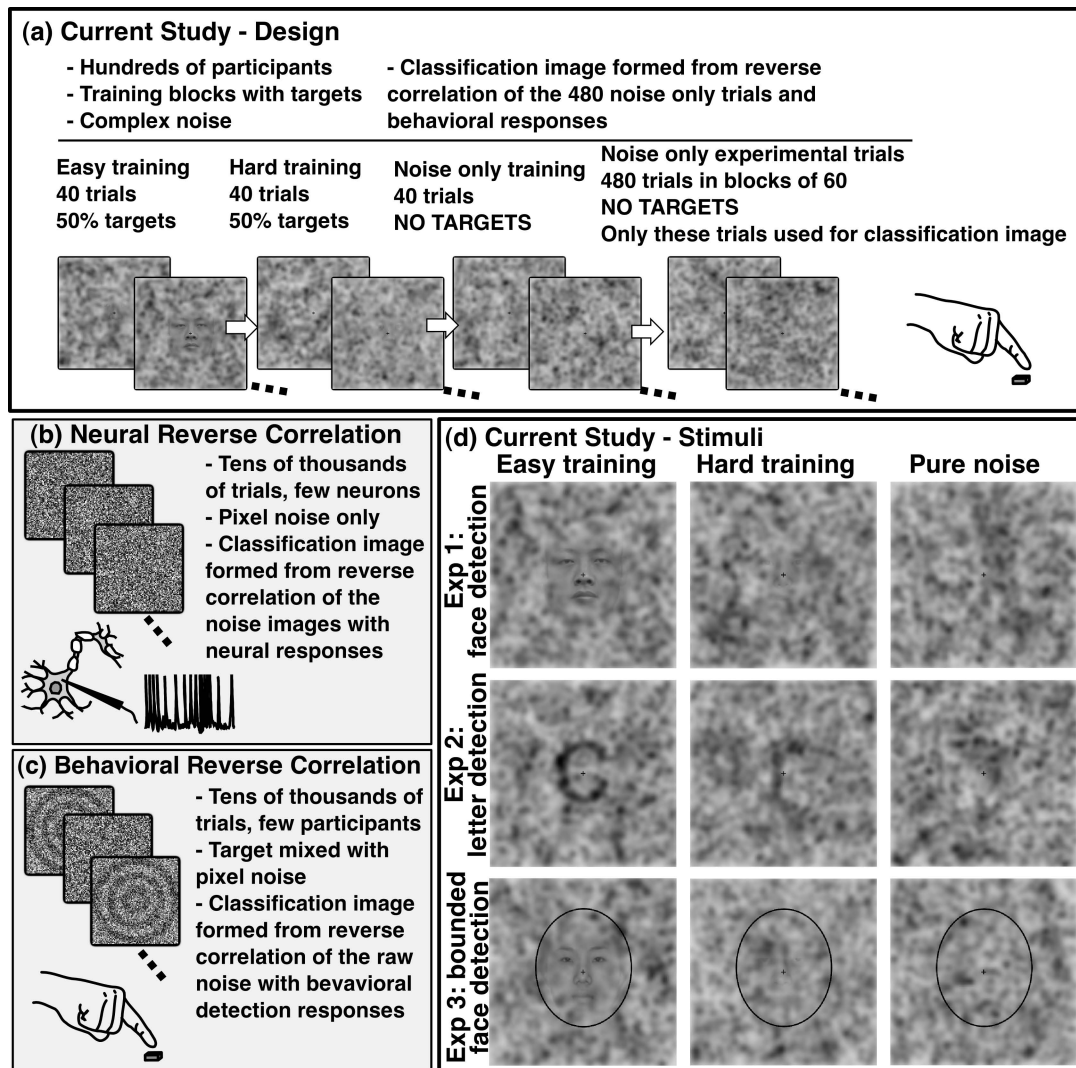


Figure 1. The top panel (a) illustrates the basic design of the experiments. Training started with 40 easy trials, 50% of which contained actual targets. The second block of 40 trials was similar, but the target faces or letters were more difficult to detect. Unknown to observers, the final 40 training trials used only pure noise images. Following training, observers responded to 480 pure noise images presented in blocks of 60. The panels (c) and (d) on the bottom left illustrate aspects of typical neural and behavioral reverse correlation experiments. The bottom right (d) shows stimuli presented to observers during Experiment 1 (top row), Experiment 2 (middle row), and Experiment 3 (bottom row). The same set of pure noise images were used for all three experiments. Target stimuli for training used 20 different pictures of male and female Asian faces in the face detection experiments, or the letters “s”, “o”, “r”, “u”, “a”, “c”, “e”, “m”, and “n” in the letter detection experiment. The faces in the hard training images may be hard to see at the size reproduced here but were clearly visible in the experiment, as evidenced by the accuracy of observers.

keep observers engaged throughout the experiment, and to reduce boredom and fatigue, we used both a smaller number of trials (480) and complex noise images. The use of complex noise images has been shown to produce reliable results and to reduce the number of trials required for behavioral reverse correlation (Hansen et al 2010). By using complex noise images that lend themselves to many interpretations, it is more plausible that a particular image might contain a target. Nevertheless, even complex noise images cannot produce a reliable CI from only 480 observations.

To gain the necessary statistical power, we analyzed the results over a large group of observers. This allowed us to statically generalize the resulting CI across observers rather than relying solely on within-observer analyses. However, it is possible that each observer might base their detection responses upon a unique, observer-specific spatial pattern. To the extent that this occurs, this introduces additional noise into the results, reducing reliability of the CI. Furthermore, if each observer viewed a different small set of noise images, the resultant CI would include sampling noise across observers and sampling noise across different noise images (i.e., both observer and noise image would be random factors). To account for this, in our paradigm we used the same set of 480 noise images for all of the observers instead of using a completely different noise image every trial. In other words, we increased reliability by using a repeated measures design. However, this introduced the issue that our results, while reliable, may be driven by the particular sample of noise images. Therefore, we also counterbalanced the orientation of the 480 noise images across different groups of observers to ensure that the particular sample of noise images did induce particular spatial attributes of the CI, such as horizontal or vertical asymmetries. This is an important control given that several variables of interest were spatial distribution analyses of the classification images (e.g., central versus peripheral or left versus right). A lateral bias that changed with the orientation of the noise images would indicate that the lateral bias was inherent in the noise images and, accordingly, that illusory detection was strongly influenced by bottom-up attributes.

2 Materials and methods

2.1 Observers

In total, 699 undergraduate students participated across three experiments: 229 in Experiment 1, 211 in Experiment 2, and 259 in Experiment 3. All observers participated for course credit, gave informed consent, and were debriefed as to the nature of the experiment. The study was approved by the University of California, San Diego (UCSD) institutional review board.

2.2 Stimuli

Noise images for the experiment were created by combining dark blobs at random spatial positions. The randomly positioned blobs were two-dimensional Gaussian distributions with three different spatial standard deviations, resulting in three different blob sizes. Furthermore, the number of randomly positioned blobs varied inversely with their size. These three spatial scales were combined to create 480 different noise images that were 480×480 pixels in size. As viewed, these images subtended 14 deg of visual angle horizontally and vertically. The same 480 noise images were used for all observers in all experiments except for the addition of an overlaid oval in Experiment 3. An additional 120 noise images were created for training. We previously used these same stimuli in several fMRI studies with Chinese observers (Li et al 2009, 2010; Liu et al 2010, 2011; Zhang et al 2008).

Every image on every trial contained a black fixation cross in the middle of the image. Target stimuli for face training were created by overlaying each noise image with a true face image that was centered. In total, there were 20 different true face images (10 male, 10 female) that were of approximately the same size. Because we sought illusory detection of any face, we could have used a mix of different size faces, although doing so would have reduced the reliability of the CI (i.e., we induced consistency in terms of the size and location of the targets but not the specific form of the targets). These 20 faces were all Asian in appearance, which was a reasonable constraint considering that the largest segment of the UCSD undergraduate population is Asian. To promote a bias for the center of the display, the image mixing proportion between the true images and the noise images was a Gaussian

distribution centered on the image, thus causing the edges of the face to fade into the noisy background. Two different mixing proportions were used for the first (easy) and second (hard) blocks of training, such that the true images were either easily seen or barely discernible. For letter training in Experiment 2 noise images were combined with true images of the letters “s”, “o”, “r”, “u”, “a”, “c”, “e”, “m”, and “n”. In pilot testing we found that letters with only straight segments were too easy to detect (i.e., illusory detection of straight letters seemed unlikely), most likely because the noise images contained no straight segments. Thus, we chose this group of letters because they contained curved line segments, which allowed the impression that the pure noise images might contain one of these letters. Participants were not informed that these were the only letters that would appear. Letter stimuli were created using the Arial typeface and blurred with a 7 pixel Gaussian blur to remove hard edges. In the bounded images used in Experiment 3 an oval roughly the size of the training faces was added to all images in the experiment. Examples are shown in [Figure 1](#).

2.3 Procedure

Observers were instructed to detect faces (Experiments 1 and 3) or letters (Experiment 2). For example, observers who detected faces with the left key were told:

You will see a number of images. The task is to press the left arrow if you think there could be a face, or the right arrow if not. The task will start easy, but will get very hard. You may feel like you are guessing but other research suggests you can be quite accurate. Focus on the fixation in between images. You should respond with the presence of a face about half the time. The images will only stay on screen briefly so respond quickly.

Observers completed three training blocks of 40 trials each, followed by six experimental blocks of 80 trials that contained only pure noise images. Between blocks, observers rested for at least 15 s. The trials within each block were presented in a random order for each observer. For the first two blocks of training half of the trials presented a true image in noise (face or letter, as appropriate), while the other half presented pure noise images. The first training block presented easy to detect targets, the second block presented hard to detect targets, and the third block presented only pure noise images. Between the training blocks they were told that the task would get harder. Before the experimental blocks they were told the difficulty of the task would be maintained. Before both the pure noise training block and the first experimental block they were encouraged to respond with detection responses about half of the time.

Each trial began with a fixation cross presented for 200 ms followed by an image that remained on the screen until a response was given or 600 ms elapsed, which was immediately followed by the fixation cross of the next trial. Observers made responses by pressing labeled keys to indicate detection or nondetection. No feedback was given. Across observers, image orientation was counterbalanced by presenting the 480 noise images either in their original orientation or 180° rotated. This manipulation demonstrated that the observed left visual field bias for face detection was not an artifact of the particular 480 pure noise images. Experiment 3 included an additional left/right control by counterbalancing response key across observers. Furthermore, in Experiment 3 all of the images (training and noise) were horizontally flipped as compared with their presentation in Experiments 1 and 2 so that the right side of the image was mapped to the left side of the image. Unlike rotation, which would invert any face-like pattern that appeared in a particular noise image, horizontal mirroring would preserve any face-like pattern but would place it in the opposite visual field.

3 Results

A few observers guessed the nature of the experiment and never detected targets during experimental trials. However, during postexperiment questioning, the vast majority of

observers indicated that they were convinced that targets appeared on some proportion of experimental trials. Across all observers the proportions of trials for which participants indicated illusory detections were 32%, 36%, and 36% for Experiments 1, 2, and 3, respectively (s.e.m = .013, .012, and .012). Observers whose hit rate for the first training block with easy targets was less than 40% or who had fewer than five illusory detections during the experimental trials were eliminated from further analysis. These criteria eliminated 24.5% of participants from Experiment 1 (173 remained, 79 with the original image orientation and 94 with the rotated), 11.0% from Experiment 2 (188 remained, 86 with the original image orientation and 102 with the rotated), and 22.4% from Experiment 3 (201 remained, 53 with mirrored images and the original response keys, 52 with mirrored images and the swapped response keys, 43 with mirrored and rotated images and the original response keys, and 53 with mirrored and rotated images and the swapped response keys). Response statistics for the remaining observers are shown in [Table 1](#).

Table 1. Mean response measures for the observers that met the criteria for inclusion.

Block	Hit rate	False alarm rate	d'
<i>Easy training</i>			
Experiment 1: face detection	0.97 (0.004)	0.05 (0.007)	3.44 (0.042)
Experiment 2: letter detection	0.91 (0.006)	0.07 (0.006)	2.94 (0.045)
Experiment 3: bounded face detection	0.96 (0.004)	0.03 (0.003)	3.48 (0.033)
<i>Hard training</i>			
Experiment 1: face detection	0.96 (0.004)	0.06 (0.007)	3.33 (0.042)
Experiment 2: letter detection	0.89 (0.007)	0.10 (0.008)	2.74 (0.050)
Experiment 3: bounded face detection	0.96 (0.003)	0.04 (0.004)	3.45 (0.033)
<i>Noise-only training</i>			
Experiment 1: face detection	na	0.29 (0.011)	na
Experiment 2: letter detection	na	0.32 (0.011)	na
Experiment 3: bounded face detection	na	0.31 (0.011)	na
<i>Noise-only experimental</i>			
Experiment 1: face detection	na	0.39 (0.011)	na
Experiment 2: letter detection	na	0.38 (0.011)	na
Experiment 3: bounded face detection	na	0.41 (0.010)	na

Note: Standard error of the means are given in parentheses; na = nonapplicable.

Because our experimental method is new we used two methods to analyze the results. Both reached the same conclusions regarding the spatial distribution of correlation values across the classification images. In the first analysis we created a separate CI for each observer in the traditional manner and compared the median correlations between different image regions using a repeated measures ANOVA. Because each observer rated only 480 images, these median correlation values based on individual CIs were small. However, because these values were analyzed across individuals, regional differences were reliable. This method of analysis allowed us to make claims regarding the population of observers, and it also allowed us to evaluate the role of between-observer control factors such as image orientation. However, because this method was based on median correlations within a spatial region, it did not allow analyses of the distribution of correlation values across pixels within a spatial region. To analyze these distributions, a second analysis method evaluated a single between-observer CI for each experiment by correlating the proportion of detection responses for

each image with the luminance of each pixel. The significance of this CI and measurements of it were then evaluated using null sampling distributions that were generated by repeatedly shuffling observer responses (i.e., Monte Carlo sampling).

3.1 Repeated measures analyses of individual CIs

In the first analysis method a CI was created for each observer by correlating their detection/nondetection response with the luminance (i.e., gray scale value) of each pixel across the 480 noise images viewed during the experimental portion of the study. The spatial distribution of CIs for each experiment were analyzed by finding the median pixel correlation within different spatial regions of interest for each observer's CI. These median correlation values were then analyzed using an ANOVA with observer as a random factor. Control factors such as image rotation and response key were included in the analyses. Correlation values were largest for the central regions, and so all analyses were performed on the center ninth of the CIs, based on an evenly spaced three-by-three grid over the entire image.

First, we compared the left and right halves of the center region (see insets at the top of Figure 2). Dark areas on the left of images were more correlated with detections than the right for both Experiment 1, $F(1, 171) = 4.19$, $p = .042$, $\eta^2 = .024$, and Experiment 3, $F(4, 197) = 6.88$, $p = .009$, $\eta^2 = .034$, but not for Experiment 2, $F(1, 186) < 1$. Next, the top and bottom halves of the center region were compared, revealing no differences for Experiment 1, $F(1, 171) < 1$, or Experiment 2, $F(1, 186) < 1$, but a slight bias toward the top for Experiment 3, $F(4, 197) = 3.91$, $p = .049$, $\eta^2 = .019$.

To provide a more fine-grained spatial analysis, the center region, which resulted from a three-by-three grid over the whole image, was itself further divided into three rows by three columns (see inset of Figure 2, upper right). For Experiment 1 the only significant effect was an effect of column, $F(2, 170) = 4.21$, $p = .016$, $\eta^2 = .047$. The pattern over columns was middle, to left, to right, and post hoc tests, $\alpha_{\text{adj}} = .017$, showed that the middle column was more correlated than the right, $t(172) = 2.98$, $p = .003$. For Experiment 2 there were significant interactions between row and column, $F(4, 183) = 5.67$, $p < .001$, $\eta^2 = .110$, as well as main effects of row, $F(2, 185) = 14.77$, $p < .001$, $\eta^2 = .13$, and column, $F(2, 185) = 17.69$, $p < .001$, $\eta^2 = .161$. Post hoc corrected t -tests assessing the interaction, $\alpha_{\text{adj}} = .0014$, found these effects were largely driven by the center region, which had much larger correlations than all other regions, $t(187) \geq 4.64$, all p -values $< .001$. The upper-middle, middle-left, middle-right, lower-left, and lower-middle regions were additionally more correlated than the lower-right region, $t(187) \geq 3.39$, all p -values $\leq .001$. For Experiment 3 there was again a row-by-column interaction, $F(4, 194) = 3.93$, $p = .004$, $\eta^2 = .075$, as well as main effects of row, $F(2, 196) = 4.32$, $p = .015$, $\eta^2 = .042$, and column, $F(2, 196) = 9.53$, $p < .001$, $\eta^2 = .089$. In contrast to Experiment 2, corrected comparisons showed that pixels in the upper-left region had stronger correlations than the middle-left, middle-right, and lower-right regions, $t \geq 3.39$, p -values $\leq .001$, and that the center and lower middle had higher correlations than the middle right and lower right, $t(200) \geq 3.26$, p -values $\leq .001$.

Next, we report the results of the control variables. These control variables were counterbalanced across observers, and the focus of these analyses was on interactions between the control variables and the spatial characteristics of the CIs. There were no interactions with image orientation for Experiment 1, F -ratios < 1.45 , p -values $> .238$. In Experiment 2, however, image orientation interacted with left/right half, $F(1, 186) = 17.94$, $p < .001$, $\eta^2 = .088$, and with top/bottom half, $F(1, 186) = 7.55$, $p < .001$, $\eta^2 = .039$. Additionally, there was an interaction with column for the three-by-three analysis, $F(2, 185) = 9.32$, $p < .001$, $\eta^2 = .092$. The interactions of image orientation with halves are shown along the left and middle columns of Figure 2. For the images shown in their original orientation, the right was significantly more correlated than the left, $t(85) = 2.26$, $p = .026$, and there was a marginal

difference such that the bottom was more correlated than the top, $t(85) = 1.93$, $p = .056$. Both of these effects reversed when the images were rotated: left more correlated than right, $t(101) = 3.82$, $p \leq .001$; top more correlated than bottom, $t(101) = 1.99$, $p = .049$. Examination of the column-by-orientation interaction, $\alpha_{\text{adj}} = .017$, showed that for the original orientation images there were differences between the center columns and each side: both $t(85) \geq 3.22$, both p -values $\leq .002$. But for the rotated images both the center and left regions had stronger correlations than the right: both $t(101) \geq 4.158$, both p -values $< .001$. In Experiment 3 there were no significant interactions with image rotation or response key for any of the analyses, F -ratios ≤ 1.34 , p -values $\geq .256$.

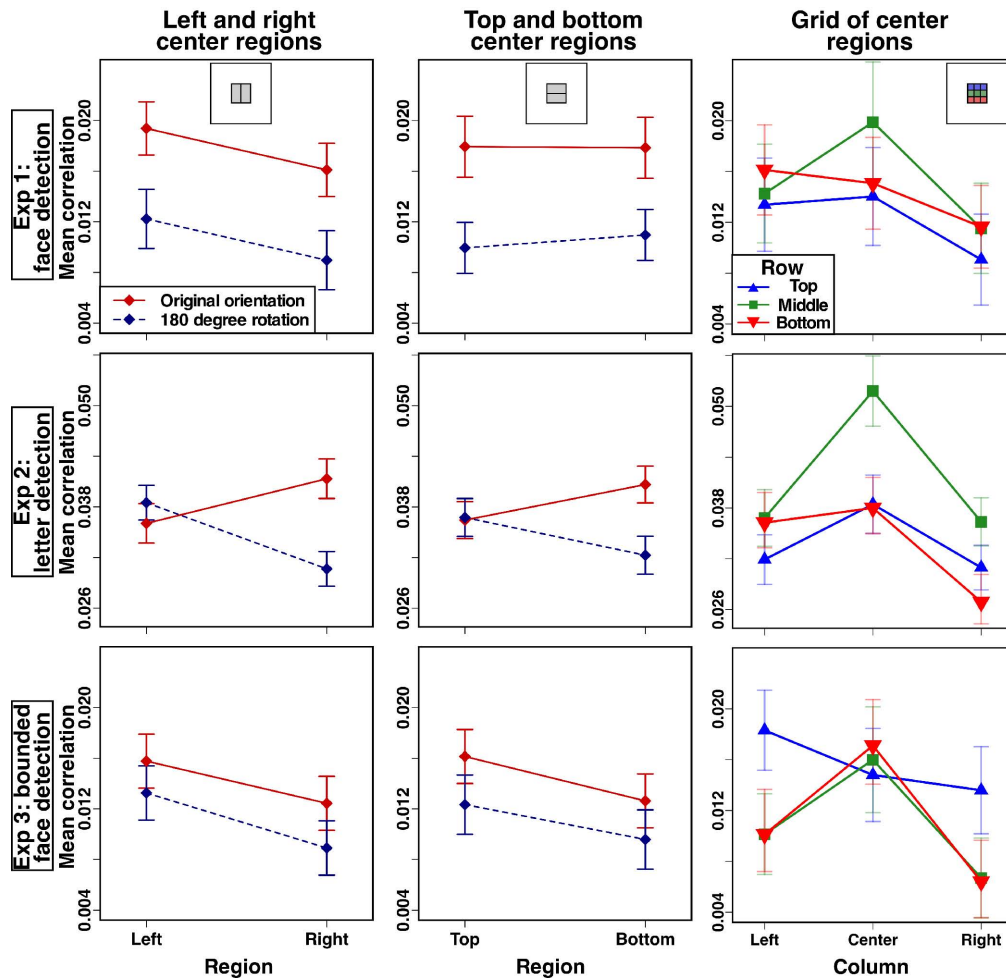


Figure 2. Repeated measures analyses of individual classification images (CIs) based on the median correlations within different regions of each observer's CI. The left column compares the left and right halves of the center region of the image as indicated by the inset at the top of the column. Between observers the 480 noise images were rotated, and the results of this control variable are shown, which produced an interaction effect only for illusory letter detection. Error bars are ± 1 standard error of the mean difference between the median left and median right correlation for each observer. The center column shows the correlations in the top and bottom center regions of the image in a similar manner. Error bars are ± 1 standard error of the mean difference between the median top and median bottom correlation for each observer. The right column compares correlation values in the three-by-three grid shown in the inset. Error bars are ± 1 standard error of the mean median correlation.

3.2 Monte Carlo analyses of between-observer CIs

A second set of analyses was based on between-observer CIs rather than on the individual CIs. A between-observer CI was calculated based on the correlation between pixel luminance (i.e., gray scale value) and the proportion of observers who gave detection responses to an image containing that pixel. The detection proportions for each image were computed as the number of detection responses divided by the sum of the number of detection and no-detection responses (i.e., excluding no-response trials). Including no-response trials did not change the results. It is worth noting that these between-observer CIs were nearly identical in appearance to the average of the individual CIs. Null hypothesis distributions were determined separately for each pixel location. Thus, if there were a systematic bias in the noise images (i.e., a tendency for dark areas to be on the left), it would be reflected in the null sampling distributions. These analyses allowed a fine-grained spatial evaluation by examining clusters of pixels within the CIs.

The expected distribution of correlation values at each pixel location from random responding was determined by running 5,000 Monte Carlo simulations for each experiment. For each Monte Carlo simulation the mapping between the 480 detection probabilities and the 480 images was randomly reshuffled. The resulting means and standard deviations at each pixel location were used to calculate z -scores for each observed pixel correlation, which are displayed in the right column of [Figure 3](#). To verify that the correlation values were normally distributed, we performed a Lilliefors test (i.e., a Kolmogorov-Smirnov test with unknown parameter values) separately for each of the 230,400 pixels of each experiment. Setting $\alpha = .05$, 5% of the pixels produced correlation distributions that differed significantly from a normal distribution (i.e., the proportion of rejections was exactly as expected based on the chosen type-I error rate).

Next, the distribution of correlation values was assessed by counting the number of pixels in the top versus bottom and in the left versus right that exceeded a chosen correlation threshold. Correlation thresholds were defined by the proportion of the largest *absolute* correlations; a rank-ordered list of all the absolute correlations for the entire image was formed from largest to smallest, and different proportions of this list were included in the analyses. As shown in [Figure 4](#), the pixel counts for one region were subtracted from the other region to yield a measure of relative spatial bias. The black lines in the figures show the observed biases, and the dashed red lines show 95% confidence intervals for this measure based on the Monte Carlo simulations. Similar to the repeated measures analyses, these analyses were carried out only for the middle region of the display, as shown by the insets of the figure. As seen in the left column of [Figure 4](#), face detection revealed a left bias (both Experiment 1 and Experiment 3), replicating the ANOVA analyses. As seen in the right column of [Figure 4](#), only bounded face detection (Experiment 3) showed any reliable vertical bias, with the most reliable correlations found in the top half.

Beyond assessing different spatial regions, the between-observer CIs were used to assess the nature of pixel clusters (i.e., correlations between pixels). To some degree, regions of correlated pixels were expected because the noise images contained Gaussian blobs rather than independent pixel noise. However, because the noise images contained Gaussian blobs of different sizes, we could evaluate biases for pixels of different cluster sizes across the three experiments. Regions were defined as groups of contiguous pixels with the absolute value of correlations above a threshold proportion. The number of contiguous regions, mean size of regions, and standard deviation of region sizes were computed for varying threshold proportions for the observed between-observer CIs and the Monte Carlo CIs. The results are plotted in [Figure 5](#). For small threshold proportions the observed data of all three experiments revealed fewer clusters than expected by chance and that those clusters were larger and more

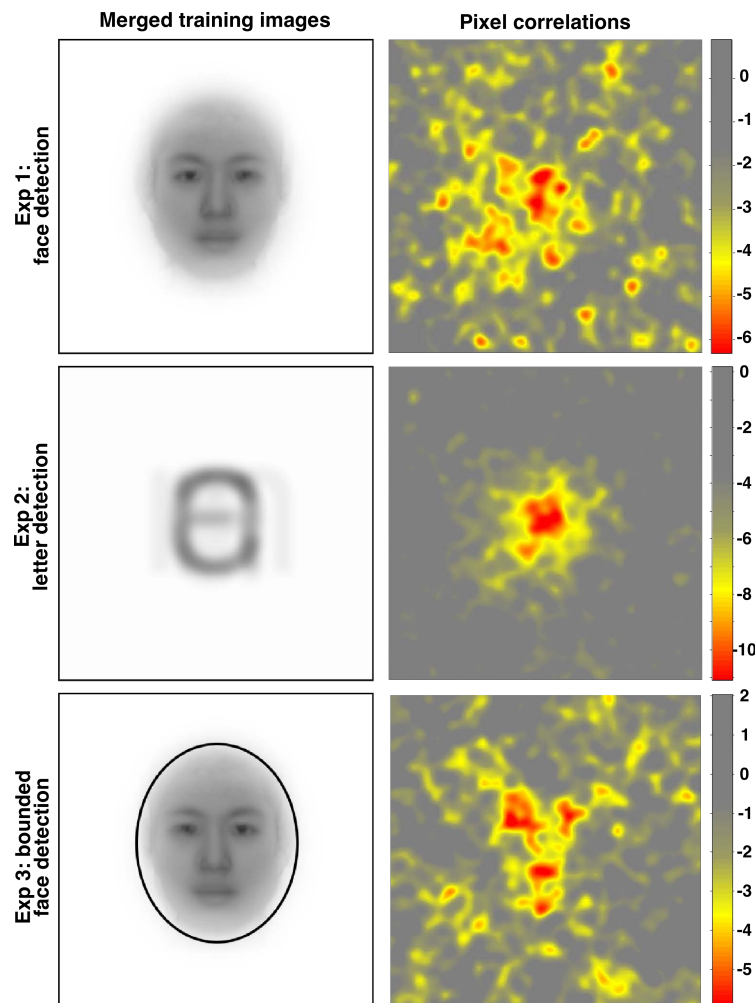


Figure 3. Left: the average of the 40 true images viewed during training. Right: the between-observer classification images (CIs) for Experiment 1 (top, $N = 173$), Experiment 2 (middle, $N = 188$), and Experiment 3 (bottom, $N = 201$). Each pixel's color is defined by its z -score as calculated from a null distribution generated by 5,000 Monte Carlo simulations. The z -score color maps are shown to the right of each CI. A Šidák correction for $\alpha = .05$ places the z -cutoff at -5.05 ; only negative z -scores (i.e., the yellow and red regions) reliably correlated with the proportion of detection responses, indicating that observers based their detections on dark rather than on light patches. The color mapping is scaled for each experiment so that the top and bottom 0.1% of correlations are shown with the same minimum or maximum color.

variable in size than would be expected. However, these effects were greatly magnified for letter detection, which had the smallest number of regions that were of the largest size.

4 Discussion

4.1 Summary of results

In a pure noise illusory detection paradigm Experiment 1 produced a distributed, but reliable, CI that was biased towards dark areas of the left side of the noise images. The observed correlations were significant and larger than would have occurred with random responding, and yet there was no obvious structure to the CI. Instead, the CI revealed a seemingly random pattern of peripheral dark patches with a left bias. This distributed pattern occurred despite the use of a fixation cross on all trials and centered faces during initial training. Experiment 2 used the same technique for a task that induced illusory letter detection, addressing the

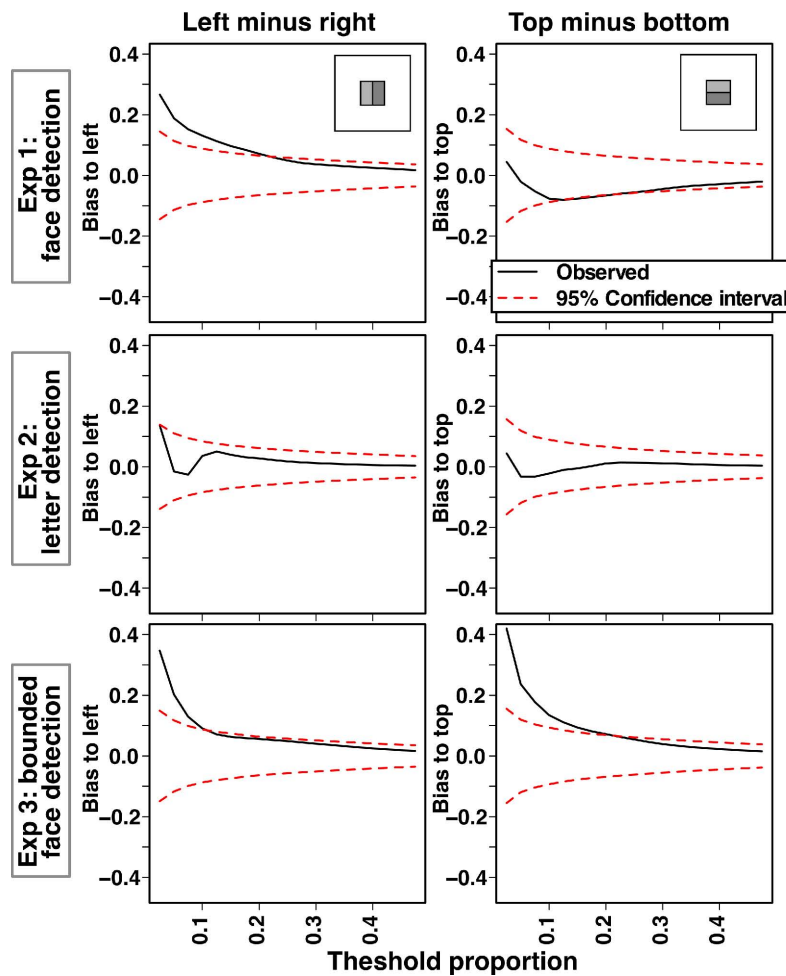


Figure 4. The distribution of correlation values within different regions of the between-observer classification images (CIs). These graphs show the difference in the number of pixels that fall into one region of the CI versus another region (e.g., left minus right) as a function of a threshold proportion that determined which pixels to include in the analysis (i.e., the threshold proportion indicates the proportion of all pixels to include in the analysis based on the absolute magnitude of correlation values). Thus, in moving from left to right, the graphs use an increasingly liberal criterion for inclusion of pixels. The black line shows this analysis based on the observed data, and the red dashed lines are 95% confidence intervals as determined by Monte Carlo sampling. The left column shows lateral bias (left minus right), and the right column shows vertical bias (top minus bottom), with these regions shown in the insets.

possibility that illusory detection of any object class would yield similar results. However, there was no consistent laterality bias, and letter detection was strongly influenced by dark patches in the center of the pure noise images (in fact, these correlations were more central than expected based on the training images, as can be seen by comparing the left and right columns of Figure 3). Furthermore, there were interactions for letter detection between the orientation of the images and the lateral biases in the images, which was not seen in the face detection experiments. Experiment 3 used the same procedure as Experiment 1 with the addition of an oval on the noise images to test if this would induce a more face-like CI by spatially constraining the location of illusory detections. The CI for Experiment 3 does have the rough appearance of a face, but still reveals a left bias. The regions corresponding to the dark areas of a face (eyes/eyebrows, nostrils, and mouth: upper left, center, and lower

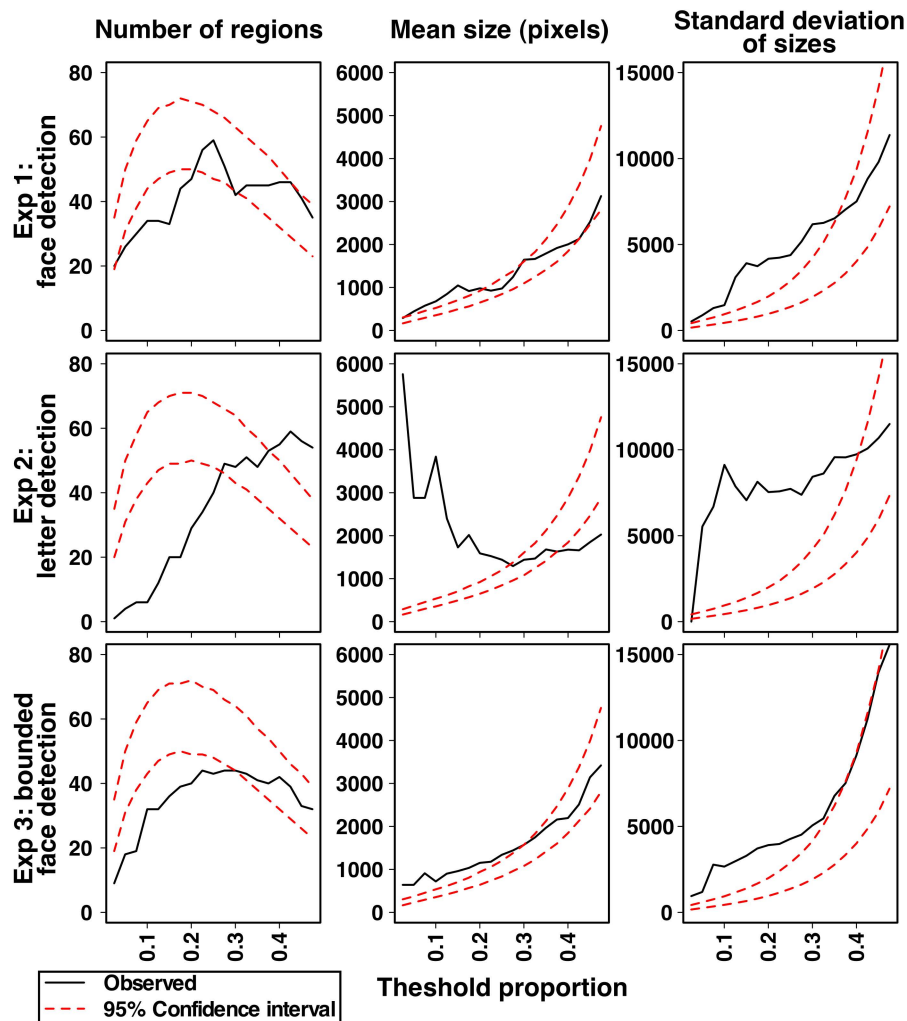


Figure 5. Monte Carlo analyses of between-observer classification images to evaluate clusters of correlated pixels. For pixel correlations greater in magnitude than the proportion threshold the number (left column), size (middle column), and standard deviation (right column) of spatially contiguous groupings of pixels are plotted. The experimentally observed values are in black, and 95% confidence intervals from the Monte Carlo sampling are plotted as red dashed lines.

middle, respectively) correlated more strongly with detection responses, with a tendency for the left to be more strongly correlated than the right.

4.2 Methodological implications and considerations

Unlike previous studies using behavioral reverse correlation, these experiments investigated illusory detection of an object class rather than a specific target exemplar. Furthermore, these experiments demonstrated that between-observer classification images can be obtained using a small number of noise images with many observers.

The use of an illusory detection paradigm (i.e., only noise images) introduces a complication in that some observers may produce much higher false alarm rates than other observers (Wenger et al 2008). To the extent that this occurs, it introduces noise in the CIs because some observers respond due to a bias to respond in general rather than respond based on a good match between the noise image and their internal template. For this reason it was important to carefully assess the reliability of the CIs for illusory detection, and so we used two different

analysis techniques to test the generality of the statistical conclusions. While there were individual differences in response rates, both analysis techniques produced qualitatively similar results: (1) illusory face detection was left lateralized despite the expectation of central targets; (2) unbounded face detection occurred in response to dark patches dispersed throughout the noise image; and (3) illusory detection of letters occurred in response to central dark patches and was more consistently driven by the bottom-up information contained in the noise images.

In interpreting these results, we must consider the training images that established the demand characteristics of the detection task. To promote reliable results, the training images appeared in the same location, were roughly the same size, and did not include the full diversity of all possible faces or letters. If training had presented faces or letters of different sizes and at different positions, then the CIs would most likely have been dispersed in all conditions (although for methods examining reverse correlation with detection at different locations, see Tjan and Nandy 2006). Rather than using training images, Gosselin and Schyns (2003) used a pure noise procedure with only verbal instructions about what to detect. Their paradigm required a few dedicated observers taking part in multiple sessions, and the target stimulus was a specific letter or the form of a mouth within the context of visible face outline rather than an object class. Their mouth detection experiment is of particular interest because it most closely corresponds to our Experiment 3 in that both studies provided the outline of the face while asking that observers detect interior components of the face. Similar to their results, Experiment 3 produced the CI that most closely resembled a face. This validated our paradigm, which was very different from the Gosselin and Schyns paradigm in many respects. This validation was important because Experiment 1 produced radically different results from other uses of reverse correlation, suggesting that observers detected faces in different regions of the pure noise images despite the use of a fixation cross and despite initial training that faces would occur only in the center of the display.

One potential concern in comparing the results of letter and face detection is that the images of the letters viewed during training were slightly smaller than the faces viewed during training (although note that, as presented in noise, their subjective sizes are comparable; see Figure 1). Even so, the CI for the letter detection task revealed a pattern that was more centralized than the *actual* letters viewed during training, whereas the CI for the face detection task of Experiment 1 revealed a pattern that was less centralized than *actual* faces viewed during training. Thus, slight differences in the size of the training objects cannot explain the difference in the spatial dispersion of the CIs.

4.3 Theoretical implications

This paradigm was originally developed to study the neural correlates of illusory detection with fMRI, which is a technique that is limited to a modest numbers of trials. The fMRI experiments identified distributed patterns of neural activation involved in top-down illusory face and letter detection (Li et al 2009, 2010; Liu et al 2010, 2011; Zhang et al 2008). By collecting data from a large number of observers, the current experiments examined behavioral results with this paradigm to ascertain the spatial patterns that tend to promote illusory detection. Thus, between the prior fMRI experiments and the current behavioral experiments we have identified both the neural and visual attributes of illusory face and letter detection.

Unlike face detection, the spatial distribution for letter detection was more tightly focused on the center. This difference is consistent with the hypothesis that top-down face processing is less constrained by spatial position, which produced the seemingly random pattern in Experiment 1 due to the superposition of separate spatially displaced face-like patterns. Unlike illusory letter detection, illusory face detection also revealed a left bias, which is

consistent with prior work with visible faces and with the anatomical location of face-specific areas of the cortex. However, this left bias in Experiment 1 could have resulted from an expectation for faces to appear on the left, or from a central form that was itself left biased (e.g., stronger emphasis on the left eye), such as observed in Experiment 3. It is notable that the observed correlations in Experiment 3 resemble the dark central features of a face (e.g., eyes/nose/mouth) but not other face features such as ears or forehead. In summary, by stripping away bottom-up visual information, we obtained support for the claim that top-down face processing is relatively unconstrained by the task demand to detect objects at the point of fixation.

We interpret this as evidence that top-down face detection is more strongly driven by form than by the location of that form. In other words, if observers responded to face-like patterns regardless of where on the screen those patterns occurred, then the summation of all the separate face detections would produce a spatially distributed pattern similar to what was observed. An alternative interpretation is that each observer consistently expected faces in a different spatial region (e.g., some observers expected faces on the left, while others expected faces on the right). The current results cannot determine whether the dispersion pattern of Experiment 1 reflects individual *observer differences* or whether it reflects individual *trial differences*. But regardless of which explanation is correct, both explanations entail detection of faces in noncentral portions of the display despite the use of central faces during training and the use of a fixation cross on all trials. These results suggest that face detection is systematically triggered by sensory information even if that information is found in unexpected locations. Thus, illusory face detection in everyday situations may be a common occurrence due to a tendency to detect faces based on a simple minimal eye/nose/mouth pattern not only within the current focus of the eyes but also for more peripheral areas of a visual scene. This does not necessarily indicate that face detection is more accurate than letter detection. Rather, because there is a high cost when failing to detect a true face (e.g., mistaking your wife for a hat, Sacks 1985), there may be a lower threshold for the information necessary to trigger face detection, even in locations where no face is expected. This results in detection of faces that do exist at the cost of occasional illusory face detection.

Both face detection experiments showed a left-side bias. This bias is consistent with behavioral studies and neural evidence. Behavioral research has established a bias toward the left side of faces (Ellis et al 1979; Gilbert and Bakan 1973; Hay 1981; Hellige et al 1984; Le Grand et al 2001; Levine and Koch-Weser 1982; Levine et al 1988; Rhodes 1985; Sergent 1984). Furthermore, studies using the “bubbles” reverse correlation technique found qualitative (Gosselin and Schyns 2001; Schyns et al 2002) and quantitative (Gosselin et al 2010) evidence for a left-side bias. Because the right hemisphere of the brain receives initial input from the left visual field, electrophysiological studies finding larger responses in the right cortical hemisphere when viewing a face are consistent with a left-side bias (Bentin et al 1996; Campanella et al 2000; Rossion et al 2003; Yovel et al 2003). This right hemisphere lateralization for faces was also confirmed with neuroimaging data (Kanwisher and Yovel 2006; Kanwisher et al 1997; Sergent et al 1992). Consistent with this mapping between the left side of a face and the right hemisphere, observers with greater left-side bias in face recognition had greater right lateralized fusiform gyrus activation when viewing faces (Yovel et al 2008).

On average, illusory letter detection showed no lateral bias even though a right-side bias might be expected due to the left hemisphere lateralization of language and reading (Binder et al 1996; Damasio and Geschwind 1984; Geschwind and Levitsky 1968; McCandliss et al 2003) and perceptual biases for the right visual field while reading and detecting letters

(Heron 1957; Mishkin and Forgy 1952; Rayner et al 1980; Robertshaw and Sheldon 1976). However, these prior studies used clearly visible letters, whereas our experiments did not include any consistent bottom-up information and created strong expectations that letters would appear only in the center. Consistent with our results, other studies using reverse correlation to study letter detection similarly failed to find spatial biases (Fiset et al 2009; Gosselin and Schyns 2003).

Compared with illusory face detection, illusory letter detection produced stronger correlations and interacted with the orientation of the specific stimuli (i.e., for letter detection the images produced a right-side bias in their original orientation but a left-side bias when rotated). This suggests that letter detection is more strongly driven by the bottom-up information contained in the noise images and so there was greater reliability between observers in determining which particular noise images did, or did not, appear to contain a letter. Furthermore, the central pattern for letters suggests that illusory letter detection is more spatially localized (e.g., the expectation that letters will be well foveated), which may have contributed to the greater between-observer reliability for letters.

Finally, we consider the relation between these results and brain imaging studies of face perception. Our results provide an important missing piece of the puzzle when interpreting neural studies of face processing. In conjunction with the currently reported behavioral results, this suggests that face processing has a frequently engaged top-down component whereby the brain creates a face interpretation based on a minimal eyes/nose/mouth pattern, regardless of visual location. These results suggest that the areas of the brain involved in face perception may be wired to automatically identify faces across the entire visual scene to a greater extent than letters.

5 Conclusions

Using a paradigm previously developed to investigate the neural correlates of illusory detection, our experiments investigated the spatial patterns that promote illusory detection of faces and letters. Our results suggest that a minimal face-like pattern is sought when detecting faces (Experiment 3) and that, in the absence of face contour illusory, face detection can occur in peripheral to the expected central location (Experiment 1). In contrast, letter detection produced larger correlation values (i.e., it occurred more consistently in response to particular intensities); and, unlike faces, letter detection revealed a strong central bias (Experiment 2). Finally, unlike letters, which did not on average produce a lateral bias, face detection was biased towards the left side of the display regardless of whether an oval was included to constrain face detection. These results suggest that face detection is more heavily top down and less constrained by task expectations (i.e., training that faces appear only in the center) as compared with letter detection.

Acknowledgements. This research was supported by NSF Grant BCS-0843773, NSF IGERT Grant DGE-0333451 to GW Cottrell/VR de Sa, and the National Natural Science Foundation of China under (NSFC) Grant No. 60910006, 31028010, 30970771. We thank Rich Shiffrin, Hal Pashler, Piotr Winkielman, and Rosie Cowell for helpful suggestions.

References

- Ahumada Jr A, 1996 "Perceptual classification images from Vernier acuity masked by noise" *Perception* **26** 1831–1840 ◀
- Andrews T J, Schluppeck D, Homfray D, Matthews P, Blakemore C, 2002 "Activity in the Fusiform Gyrus predicts conscious perception of Rubin's vase-face illusion" *NeuroImage* **17** 890–901 [doi:10.1006/nimg.2002.1243](https://doi.org/10.1006/nimg.2002.1243) ◀

- Bentin S, Allison T, Puce A, Perez E, McCarthy G, 1996 “Electrophysiological studies of face perception in humans” *Journal of Cognitive Neuroscience* **8** 551–565 doi:10.1162/jocn.1996.8.6.551 ◀
- Binder J, Swanson S, Hammeke T, Morris G, Mueller W, Fischer M, Benbadis S, Frost J A, Rao S M, Houghton V M, 1996 “Determination of language dominance using functional MRI” *Neurology* **46** 978–984 ◀
- Campanella S, Hanoteau C, Dépy D, Rossion B, Bruyer R, Crommelinck M, Guérit J M, 2000 “Right N170 modulation in a face discrimination task: an account for categorical perception of familiar faces” *Psychophysiology* **37** 796–806 doi:10.1111/1469-8986.3760796 ◀
- Damasio A R, Geschwind N, 1984 “The neural basis of language” *Annual Review of Neuroscience* **7** 127–147 doi:10.1146/annurev.ne.07.030184.001015 ◀
- Dayan P, Abbott L F, 2001 *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (Cambridge, MA: MIT Press) ◀
- Ellis H D, Shepherd J W, Davies G M, 1979 “Identification of familiar and unfamiliar faces from internal and external features: some implications for theories of face recognition” *Perception* **8** 431–439 doi:10.1068/p080431 ◀
- Esterman M, Yantis S, 2010 “Perceptual expectation evokes category-selective cortical activity” *Cerebral Cortex* **20** 1245–1253 doi:10.1093/cercor/bhp188 ◀
- Fairhall S L, Ishai A, 2007 “Effective connectivity within the distributed cortical network for face perception” *Cerebral Cortex* **17** 2400–2406 doi:10.1093/cercor/bhl148 ◀
- Fiset D, Blais C, Arguin M, Tadros K, Ethier-Majcher C, Bub D, Gosselin F, 2009 “The spatio-temporal dynamics of visual letter recognition” *Cognitive Neuropsychology* **26** 23–35 doi:10.1080/02643290802421160 ◀
- Geschwind N, Levitsky W, 1968 “Human brain: left-right asymmetries in temporal speech region” *Science* **161** 186–187 doi:10.1126/science.161.3837.186 ◀
- Gilbert C, Bakan P, 1973 “Visual asymmetry in perception of faces” *Neuropsychologia* **11** 355–362 doi:10.1016/0028-3932(73)90049-3 ◀
- Gold J M, Murray R F, Bennett P J, Sekuler A B, 2000 “Deriving behavioural receptive fields for visually completed contours” *Current Biology* **10** 663–666 doi:10.1016/S0960-9822(00)00523-6 ◀
- Gosselin F, Schyns P G, 2001 “Bubbles: a technique to reveal the use of information in recognition tasks” *Vision Research* **41** 2261–2271 doi:10.1016/S0042-6989(01)00097-9 ◀
- Gosselin F, Schyns P G, 2003 “Superstitious perceptions reveal properties of internal representations” *Psychological Science* **14** 505–509 doi:10.1111/1467-9280.03452 ◀
- Gosselin F, Spezio M L, Tranel D, Adolphs R, 2010 “Asymmetrical use of eye information from faces following unilateral amygdala damage” *Social Cognitive and Affective Neuroscience* **6** 1–8 ◀
- Hansen B C, Thompson B, Hess R F, Ellemberg D, 2010 “Extracting the internal representation of faces from human brain activity: An analogue to reverse correlation” *NeuroImage* **51** 373–390 doi:10.1016/j.neuroimage.2010.02.021 ◀
- Hasson U, Hendler T, Ben-Bashat D, Malach R, 2001 “Vase or face? A neural correlate of shape-selective grouping processes in the human brain” *Journal of Cognitive Neuroscience* **13** 744–753 doi:10.1162/08989290152541412 ◀
- Hay D C, 1981 “Asymmetries in face processing: evidence for a right hemisphere perceptual advantage” *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology* **33** 267–274 ◀
- Hellige J B, Corwin W H, Jonsson J E, 1984 “Effects of perceptual quality on the processing of human faces presented to the left and right cerebral hemispheres” *Journal of Experimental Psychology: Human Perception and Performance* **10** 90–107 doi:10.1037/0096-1523.10.1.90 ◀
- Heron W, 1957 “Perception as a function of retinal locus and attention” *The American Journal of Psychology* **70** 38–48 doi:10.2307/1419227 ◀
- Ishai A, 2008 “Let’s face it: It’s a cortical network” *NeuroImage* **40** 415–419 doi:10.1016/j.neuroimage.2007.10.040 ◀
- Jones J P, Palmer L A, 1987 “The two-dimensional spatial structure of simple receptive fields in cat striate cortex” *Journal of Neurophysiology* **58** 1187–1211 ◀
- Kanwisher N, Yovel G, 2006 “The fusiform face area: a cortical region specialized for the perception of faces” *Philosophical Transactions of the Royal Society B: Biological Sciences* **361** 2109–2128 doi:10.1098/rstb.2006.1934 ◀
- Kanwisher N, McDermott J, Chun M M, 1997 “The fusiform face area: a module in human extrastriate cortex specialized for face perception” *Journal of Neuroscience* **17** 4302–4311 ◀
- Kontsevich L L, Tyler C W, 2004 “What makes Mona Lisa smile?” *Vision Research* **44** 1493–1498 doi:10.1016/j.visres.2003.11.027 ◀

-
- Le Grand R, Mondloch C J, Maurer D, Brent H P, 2001 "Neuroperception: Early visual experience and face processing" *Nature* **410** 890–890 doi:10.1038/35073749 ◀
- Levine S C, Koch-Weser M P, 1982 "Right hemisphere superiority in the recognition of famous faces" *Brain and Cognition* **1** 10–22 doi:10.1016/0278-2626(82)90003-3 ◀
- Levine S C, Banich M T, Koch-Weser M P, 1988 "Face recognition: a general or specific right hemisphere capacity?" *Brain and Cognition* **8** 303–325 doi:10.1016/0278-2626(88)90057-7 ◀
- Lewis M B, Edmonds A J, 2003 "Face detection: Mapping human performance" *Perception* **32** 903–920 doi:10.1068/p5007 ◀
- Li J, Liu J, Liang J, Zhang H, Zhao J, Huber D E, Rieth C A, Lee K, Tian J, Shi G, 2009 "A distributed neural system for top-down face processing" *Neuroscience Letters* **451** 6–10 doi:10.1016/j.neulet.2008.12.039 ◀
- Li J, Liu J, Liang J, Zhang H, Zhao J, Rieth C A, Huber D E, Li W, Shi G, Ai L, Tian J, Lee K, 2010 "Effective connectivities of cortical regions for top-down face processing: a dynamic causal modeling study" *Brain Research* **1340** 40–51 doi:10.1016/j.brainres.2010.04.044 ◀
- Lindsay R, Lea J A, Nosworthy G J, Fulford J A, Hector J, LeVan V, Seabrook C, 1991 "Biased line-ups: Sequential presentation reduces the problem" *Journal of Applied Psychology* **76** 796–802 doi:10.1037/0021-9010.76.6.796 ◀
- Liu J, Li J, Zhang H, Rieth C A, Huber D E, Li W, Lee K, Tian J, 2010 "Neural correlates of top-down letter processing" *Neuropsychologia* **48** 636–641 doi:10.1016/j.neuropsychologia.2009.10.024 ◀
- Liu J, Li J, Rieth C A, Huber D E, Tian J, Lee K, 2011 "A dynamic causal modeling analysis of the effective connectivities underlying top-down letter processing" *Neuropsychologia* ◀
- McCandliss B D, Cohen L, Dehaene S, 2003 "The visual word form area: expertise for reading in the fusiform gyrus" *Trends in Cognitive Sciences* **7** 293–299 doi:10.1016/S1364-6613(03)00134-7 ◀
- Mishkin M, Forgays D G, 1952 "Word recognition as a function of retinal locus" *Journal of Experimental Psychology* **43** 43–48 doi:10.1037/h0061361 ◀
- Movshon J, Thompson I, Tolhurst D, 1978 "Spatial summation in the receptive fields of simple cells in the cat's striate cortex" *The Journal of Physiology* **283** 53–57 ◀
- O'Craven K M, Kanwisher N, 2000 "Mental imagery of faces and places activates corresponding stimulus-specific brain regions" *Journal of Cognitive Neuroscience* **12** 1013–1023 doi:10.1162/08989290051137549 ◀
- Rayner K, Well A D, Pollatsek A, 1980 "Asymmetry of the effective visual field in reading" *Attention, Perception, & Psychophysics* **27** 537–544 ◀
- Rhodes G, 1985 "Lateralized processes in face recognition" *British Journal of Psychology* **76** 249–271 doi:10.1111/j.2044-8295.1985.tb01949.x ◀
- Ringach D L, Shapley R, 2004 "Reverse correlation in neurophysiology" *Cognitive Science* **28** 147–166 doi:10.1207/s15516709cog2802_2 ◀
- Robertshaw S, Sheldon M, 1976 "Laterality effects in judgment of the identity and position of letters: a signal detection analysis" *The Quarterly Journal of Experimental Psychology* **28** 115–121 doi:10.1080/14640747608400544 ◀
- Rossion B, Joyce C A, Cottrell G W, Tarr M J, 2003 "Early lateralization and orientation tuning for face, word, and object processing in the visual cortex" *NeuroImage* **20** 1609–1624 doi:10.1016/j.neuroimage.2003.07.010 ◀
- Sacks O W, 1985 *The Man Who Mistook His Wife for a Hat and Other Clinical Tales* (New York, NY: Touchtone) ◀
- Schyns P G, Bonnar L, Gosselin F, 2002 "Show me the features! Understanding recognition from the use of visual information" *Psychological Science* **13** 402–409 doi:10.1111/1467-9280.00472 ◀
- Sergent J, 1984 "Configural processing of faces in the left and the right cerebral hemispheres" *Journal of Experimental Psychology* **10** 554–572 ◀
- Sergent J, Ohta S, MacDonald B, 1992 "Functional neuroanatomy of face and object processing: A Positron Emission Tomography study" *Brain* **115** 15–36 doi:10.1093/brain/115.1.15 ◀
- Smith F W, Muckli L, Brennan D, Pernet C, Smith M L, Belin P, Gosselin F, Hadley D M, Cavanagh J, Schyns P G, 2008 "Classification images reveal the information sensitivity of brain voxels in fMRI" *NeuroImage* **40** 1643–1654 doi:10.1016/j.neuroimage.2008.01.029 ◀
- Smith M L, Fries P, Gosselin F, Goebel R, Schyns P G, 2009 "Inverse mapping the neuronal substrates of face categorizations" *Cerebral Cortex* **19** 2428–2438 doi:10.1093/cercor/bhn257 ◀
- Svoboda E, 2007 "Faces everywhere" *The New York Times* page D1 ◀
- Tjan B S, Nandy A S, 2006 "Classification images with uncertainty" *Journal of Vision* **6** 387–413 doi:10.1167/6.4.8 ◀

-
- Wenger M J, Copeland A M, Bittner J L, Thomas R D, 2008 “Evidence for criterion shifts in visual perceptual learning: Data and implications” *Perception & Psychophysics* **70** 1248–1273 [doi:10.3758/PP.70.7.1248](https://doi.org/10.3758/PP.70.7.1248) ◀
- Yovel G, Levy J, Grabowecky M, Paller K A, 2003 “Neural correlates of the left-visual-field superiority in face perception appear at multiple stages of face processing” *Journal of Cognitive Neuroscience* **15** 462–474 [doi:10.1162/089892903321593162](https://doi.org/10.1162/089892903321593162) ◀
- Yovel G, Tambini A, Brandman T, 2008 “The asymmetry of the fusiform face area is a stable individual characteristic that underlies the left-visual-field superiority for faces” *Neuropsychologia* **46** 3061–3068 [doi:10.1016/j.neuropsychologia.2008.06.017](https://doi.org/10.1016/j.neuropsychologia.2008.06.017) ◀
- Zhang H, Liu J, Huber D E, Rieth C A, Tian J, Lee K, 2008 “Detecting faces in pure noise images: a functional MRI study on top-down perception” *NeuroReport* **19** 229–233 [doi:10.1097/WNR.0b013e3282f49083](https://doi.org/10.1097/WNR.0b013e3282f49083) ◀