

# UC Berkeley

## UC Berkeley Previously Published Works

### Title

Individual differences in classification images of Mooney faces.

### Permalink

<https://escholarship.org/uc/item/16r9b81z>

### Journal

Journal of Vision, 22(13)

### Authors

Canas-Bajo, Teresa

Whitney, David

### Publication Date

2022-12-01

### DOI

10.1167/jov.22.13.3

Peer reviewed

# Individual differences in classification images of Mooney faces

**Teresa Canas-Bajo**

Vision Science Graduate Group, University of California,  
Berkeley, Berkeley, CA, USA



Vision Science Graduate Group, University of California,  
Berkeley, Berkeley, CA, USA

**David Whitney**

Department of Psychology, University of California,  
Berkeley, Berkeley, CA, USA



Human face recognition is robust even under conditions of extreme lighting and in situations where there is high noise and uncertainty. Mooney faces are a canonical example of this: Mooney faces are two-tone shadow-defined images that are readily and holistically recognized despite lacking easily segmented face features. Face perception in such impoverished situations—and Mooney face perception in particular—is often thought to be supported by comparing encountered faces to stored templates. Here, we used a classification image approach to measure the templates that observers use to recognize Mooney faces. Visualizing these templates reveals the regions and structures of the image that best predict individual observer recognition, and they reflect the underlying internal representation of faces. Using this approach, we tested whether there are classification images that are consistent from session to session, whether the classification images are observer-specific, and whether they allow for pattern completion of holistic representations even in the absence of an underlying signal. We found that classification images of Mooney faces were indeed non-random (i.e., consistent session from session) within each observer, but they were different between observers. This result is in line with previously proposed existence of face templates that support face recognition, and further suggests that these templates may be unique to each observer and could drive idiosyncratic individual differences in holistic face recognition. Moreover, we found classification images that reflected information within the blank regions of the original Mooney faces, suggesting that observers may fill in missing information using idiosyncratic internal information about faces.

## Introduction

Faces play a central role in our everyday life, and our visual system is in turn extremely sensitive to them. Faces convey critical social and emotional information;

they guide interactions and our everyday behavior (Costela & Woods, 2019; Jack & Schyns, 2015; Coutrot & Guyader, 2014). Reinforced by a lifelong exposure to faces, humans have developed a remarkable preference for faces over other objects from infancy (Umiltà, Simion & Valenza, 1996; Farzin, Rivera & Whitney, 2009; Frank, Vul & Johnson, 2009; Pascalis & Kelly, 2009). We quickly find faces in a scene and direct our attention and gaze toward them (Bindemann, Burton, Hooge, Jenkins & de Haan, 2005; Boucart, Lenoble, Quettelart, Szaffarczyk, Despretz, & Thorpe, 2016; Cerf, Frady & Koch, 2009; Cerf, Harel, Huth, Einhauser & Koch, 2008; Costela & Woods, 2019; Coutrot & Guyader, 2014; Crouzet, Kirchner & Thorpe, 2010; Foulsham, Cheng, Tracy, Henrich & Kingstone, 2010; Jack & Schyns, 2015; Marat, Rahman, Pellerin, Guyader & Houzet, 2013; Martin, Davis, Riesenhuber, & Thorpe, 2018; Ro, Russell & Lavie, 2001; Theeuwes & Van der Stigchel, 2006; Wolfe & Whitney, 2015; Xia, Manassi, Nakayama, Zipser & Whitney, 2020). In fact, humans can recognize faces under a wide range of conditions, including under challenging environments (Burton, Bruce & Hancock, 1999; Hole, George, Eaves & Rasek, 2002). Humans have also developed expert processing mechanisms specific to faces such as holistic processing, the perception of faces as a whole (Farah, Wilson, Drain & Tanaka, 1998; Maurer, Le Grand & Mondloch, 2002; Sergent, 1984). Faces therefore enjoy a unique and privileged position in human visual processing.

Previous work suggests that face recognition may involve comparing encountered faces to internal face representations – templates (Cavanagh, 1991; Damasio, Damasio & Van Hoesen, 1982; Sekuler & Abrams, 1968; Sergent, 1984; Simpson & Crandall, 1972; Smith & Nielsen, 1970; Valentine, 1991). Faces that best match those templates are processed more efficiently than those that do not (Rossion, 2008; Rossion & Boremanse, 2008). This matching process allows the recovery of two- and three-dimensional information

Citation: Canas-Bajo, T., & Whitney, D. (2022). Individual differences in classification images of Mooney faces. *Journal of Vision*, 22(13):3, 1–16, <https://doi.org/10.1167/jov.22.13.3>.





Figure 1. Mooney faces used in this study in upright (top) and inverted (bottom) orientation. Mooney faces are two-tone images that can be quickly recognized as faces despite lacking low-level face features. Mooney faces are processed as a whole; note that although upright Mooney faces are easily perceived as faces, inverted Mooney faces are difficult or even impossible to recognize as such. For this reason, Mooney faces are ideal stimuli to test holistic processing.

from the observed faces, even in impoverished or extreme lighting conditions (Cavanagh, 1991). Valentine (1991) defined the norm-based coding model, which proposes that faces are encoded as distance vectors to a prototype shaped by individual prior experience. Others have aimed to define the structures of these templates and proposed similar models of face recognition that occur by matching observed faces to template-like structures (e.g., Bar codes: Dakin & Watt, 2009; Integral Images: Viola & Jones, 2004; Face Recognition Units: Bruce & Young, 1986; Young & Bruce, 2011). In these models, faces can be encoded in terms of low-level features that can be parsed or easily segmented. Although these models can accommodate grayscale faces well, they may not be as effective for faces that require holistic processing. For example, they would not be as successful for two-tone shadow defined faces (Mooney faces; Figure 1) because the local features, edges, contours, and spatial frequency content of Mooney faces are not discriminative of the image as a face (Andrews & Schluppeck, 2004; Bona, Cattaneo & Silvanto, 2016; Cavanagh, 1991; Farzin et al., 2009; Latinus & Taylor, 2005; McKone, 2004; Moore & Cavanagh, 1998). In principle, template-based models are an appropriate way to analyze Mooney faces (Cavanagh, 1991), but the templates need to carry information about the holistic aspects of the face.

Previous work has shown that there are individual differences in holistic processing of both grayscale and two-tone Mooney faces: some individuals process faces more holistically than others (DeGutis, Mercado, Wilmer, & Rosenblatt, 2013; DeGutis, Wilmer, Mercado, Cohan, 2013; Gauthier, 2020; Russell, Duchaine & Nakayama, 2009; Wang, Li,

Fang, Tian & Liu, 2012). Interestingly, there are also stimulus-specific individual differences in holistic processing (Canas-Bajo & Whitney, 2020): specific faces that are processed holistically by one observer are not necessarily processed holistically by other observers. The origin of these differences remains unclear. Here, our goal was to test whether such idiosyncrasy in holistic processing is caused by underlying individual differences in observer-specific face templates.

Capturing or visualizing templates can be challenging because they are implicit representations. One of the most widely used and established techniques for measuring templates is the reverse-correlation image classification method (see Murray, 2011 for a review). This technique is a data-driven method that has been widely used in the literature to capture internal representations—the image regions and structures that are most important for detection, discrimination, or classification (Ahumada, 1996, 2002; Ahumada & Lovell, 1971; Beard & Ahumada, 1998; Mareschal, Dakin & Bex, 2006; Neri & Heeger, 2002; Ringach & Shapley, 2004; Solomon, 2002; Victor, 2005; Watson, 1998). More recently, the method has also been used to visualize face templates (Brinkman, Todorov & Dostch, 2017; Dostch & Todorov, 2012; Dotsch, Wigboldus, Langner & van Knippenberg, 2008; Éthier-Majcher, Joubert & Gosselin, 2013; Gosselin and Schyns, 2003; Karremans, Dotsch, & Corneille, 2011). In a standard reverse-correlation experiment, the stimuli consist of random noise overlaid over the same base face. The observers' task is to classify the stimuli based on some construct (e.g., trustworthiness: Éthier-Majcher, Joubert & Gosselin, 2013). Classification images are

computed by averaging the noise patterns of the stimuli that participants classified as being representative of the construct of interest. Classification images can therefore serve as a visualization of internal representations of faces (Brinkman, Todorov & Dostch, 2017; Dostch & Todorov, 2012; Gosselin & Schyns, 2003; Sekuler, Gaspar, Gold & Bennett, 2004). The reverse-correlation technique is ideal to tap into implicit representations of faces because it is a purely data-driven method: noise is generated randomly in each trial for each participant, and participants use their own internal representation of the construct without external biases (Mangini & Biederman, 2004; Todorov, Dotsch, Wigboldus, & Said, 2011). In the field of face recognition, this method has been used to understand the basis of face recognition at the group level (Mangini & Biederman, 2004; Sekuler et al., 2004) and how it is influenced by factors like emotion (Brooks & Freeman, 2018; Karremans, Dostch & Corneille, 2011; Jack, Caldara & Schyns, 2012), perceptual disorders (Brinkman, Dotsch, Zondergeld, Koevoets, Aarts & van Haren, 2019), prejudice (Dotsch, Wigboldus, & Van Knippenberg, 2013), social cognition (Brooks, Stolier & Freeman, 2018), and culture (Dotsch et al., 2008; Jack et al., 2012; Ratner, Dotsch, Wigboldus, van Knippenberg, & Amodio, 2014). For instance, Dotsch and Todorov (2012) used classification images to identify racial biases in the perception of trustworthiness in faces.

Here, we use the reverse correlation technique to investigate whether there are individual differences in classification images of Mooney faces. Mooney faces are highly impoverished two-tone black and white blobs that are readily perceived as faces despite lacking low-level face features that can be parsed in a bottom-up fashion (Figure 1; Brodski, Paasch, Helbling & Wibrall, 2015; Goold & Meng, 2016; Ke, Stella, & Whitney, 2017; Mooney, 1957; Schwiedrzik, Melloni & Schurger, 2018). That is, the image must be recognized as a face before any feature (e.g., an eye, a mouth, etc) can be identified or localized (Cavanagh, 1991; Farzin et al., 2009). Consequently, holistic processing plays a critical role in the recognition of Mooney faces. Holistic processing breaks down with inverted faces, which makes inverted Mooney faces very difficult to recognize (Andrews & Schluppeck, 2004; Bona, Cattaneo & Silvanto, 2016; Cavanagh, 1991; Farzin et al., 2009; Latinus & Taylor, 2005; McKone, 2004; Moore & Cavanagh, 1998; Sergent, 1984). For these reasons, Mooney faces are ideal stimuli to isolate holistic processing (Figure 1).

The following experiments tested several hypotheses. First, because our goal was to measure the templates that support holistic perception of Mooney faces, we expected that classification images would be consistent from day to day. Second, if previously found individual differences in holistic processing reflect underlying observer-

specific templates (Canas-Bajo & Whitney, 2020), we should find idiosyncratic classification images only for upright but not inverted Mooney faces. Last, if Mooney faces are recognized by matching the stimuli to stored templates, classification images of Mooney faces should reflect filled-in information that is missing in the original, impoverished stimuli. Altogether, our study aimed to investigate whether there are observer-specific face templates (Experiment 1) and whether they allow for pattern completion of the missing information in the original Mooney faces (Experiment 2).

## Experiment 1

### Methods

#### Participants

Seven participants (three male, four female) took part in this experiment. One participant was excluded because they did not complete all trials of the experiment. All subjects were recruited through Mechanical Turk and provided written consent forms before participation. All experimental procedures were approved by the UC Berkeley Institutional Review Board.

#### Material and design

The base faces in the reverse correlation task consisted of four unfamiliar Mooney faces created by Schwiedrzik, Melloni & Schurger (2018; face labels in the original dataset: U0034, U0129, U0159, U0393). These Mooney faces were presented both in upright and inverted orientations, generating a total of eight base faces for the reverse correlation experiment (Figure 1). In order to assess the consistency of the subjects' classification images, participants did the experiment twice in two separate days. The time interval between the two sessions was on average 3.4 days (std: 1.5 days). In each experimental session, participants completed 8 blocks of 250 trials for each base image. Therefore each participant completed 4000 trials in total: base images  $\times$  250 trials per base image  $\times$  2 sessions (2000 trials per session).

The presentation of the stimuli and the data collection was controlled using Qualtrics (<https://www.qualtrics.com>). Note that participants used different monitors, so distance to monitor could not be controlled.

#### Procedure

Stimuli were generated uniquely for each participant. In each trial, we generated a random pattern of




Figure 2. (A) To generate the stimuli for each trial, the same base image was overlaid with two identical but polarity-reversed random patterns of noise. (B) Example trial. After a blank screen, two stimuli were presented, and participants were asked to select which of two images was more face-like. (C) Classification images were generated by averaging the selected noise patterns across trials.

sinusoidal noise and its corresponding negative version (Figure 2A). The random sinusoidal noise is defined by 4092 parameters, each defining the amplitude of one truncated sinusoid ranging over two cycles (for more details on how the random noise was generated, see Dostch & Todorov, 2012).

Each trial consisted of two side-by-side stimuli: the base face with the original noise superimposed and the same base face with the polarity reversed noise superimposed (Figure 2C). Participants' task was to select which of the two images presented was more face-like by pressing the left or right button. The stimuli were presented until observers gave a response. Participants could move their eyes freely during the trial and could take a break every 250 trials. The patterns of noise were randomized for each trial separately for each observer. That is, no two subjects saw the same stimuli, and the same observer did not see the same stimuli more than once. All participants completed a demographics questionnaire followed by a question regarding their previous experience with Mooney faces.

### Data analysis

Our goal here was to investigate whether face templates support recognition of Mooney faces, and whether the templates, if present, are observer-specific. We used classification images to visualize observers' face templates. To generate the classification images, we averaged all of the observer's selected noise patterns across trials, and we superimposed the resulting noise on the base face (Figure 2C). The classification image reveals the regions or structures that observers rely on to classify an image as face-like. We generated one classification image per base face, per subject. The outcome was eight unique classification images for each observer.

To evaluate whether observers' classification images were nonrandom, we tested the day-to-day consistency of each observers' classification images for each of the eight base Mooney faces. The within-subject consistency was defined as the test-retest pixel-wise correlation between an observer's classification images in the first and second session. We then averaged all within-observer Fisher  $z$  transformed correlations for upright and inverted base faces, separately. Second, we quantified the between-observer agreement in classification images of Mooney faces. To this end, we calculated the between-observer correlations across different observers' classification images for each base face. Then we averaged all between-observer Fisher  $z$  transformed correlations for upright and inverted base faces, separately.

To test the significance of the within- and between-subject correlations, we calculated a null distribution of permuted correlations. In each iteration of the permutation, for each base face, we shuffled the responses that were given by each participant. We then calculated a classification image from these shuffled responses. This was effectively the classification images an observer would have if the participant responded randomly. Using this data, we generated a set of 1000 permuted classification images, per base face, per participant. Note that our classification images have signal included, which could artificially inflate the correlations in both within- and between-subject analyses. To control for this, we calculated null distributions of permuted correlations, which maintain all of the same signal information, but represent shuffled responses, so any inflation of the correlations will occur in the null distribution as well. In other words, we never compare correlations to zero, always to permuted shuffled null distributions. To generate the distribution of permuted null within-subject correlations, for each observer and each base face, we correlated the empirical



classification image of the first session with each of the permuted null classification images of the second session. In the same fashion, to generate the distribution of permuted null between-subject correlations, for each pair observers and each base face, we correlated an observer's one-session empirical classification image with each of the permuted null classification images of another observer's session.

Last, we quantified the individual differences of classification images of Mooney faces by comparing within- and between-observer agreement. The significance of this comparison was tested using a nonparametric permutation method: we computed the empirical difference for within- and between-correlations and compared it to the null difference for the between and within correlations for randomized responses. Across all analyses, to calculate the average within- and between-subject correlations across base faces or across observers, correlations were first

transformed from Pearson  $R$  correlations into Fisher  $z$  correlations before averaging them.

## Results

In this study, we measured classification images of Mooney faces to investigate whether there are face templates that support holistic face recognition and whether these templates were observer specific. In our experiment, participants classified which stimuli were more face-like.

The outcome per observer was a classification image for each base face tested (see [Figure 3](#) for some examples). Casual inspection hints that the classification images of upright faces seem to be visible and face-like ([Figure 3A](#)). In contrast, classification images of inverted faces appeared to be less recognizable ([Figure 3B](#)). Our hypothesis was that an observer's




Figure 3. Sample resulting classification images of Mooney faces in [Experiment 1](#) for two subjects. (A) Top row: upright Mooney faces (in red). The Mooney face Classification images are recognizable (albeit noisy and ghostly). (B) Bottom row: inverted Mooney faces (in blue). Inverted Mooney face classification images did not appear to be very obvious. See main text for details.

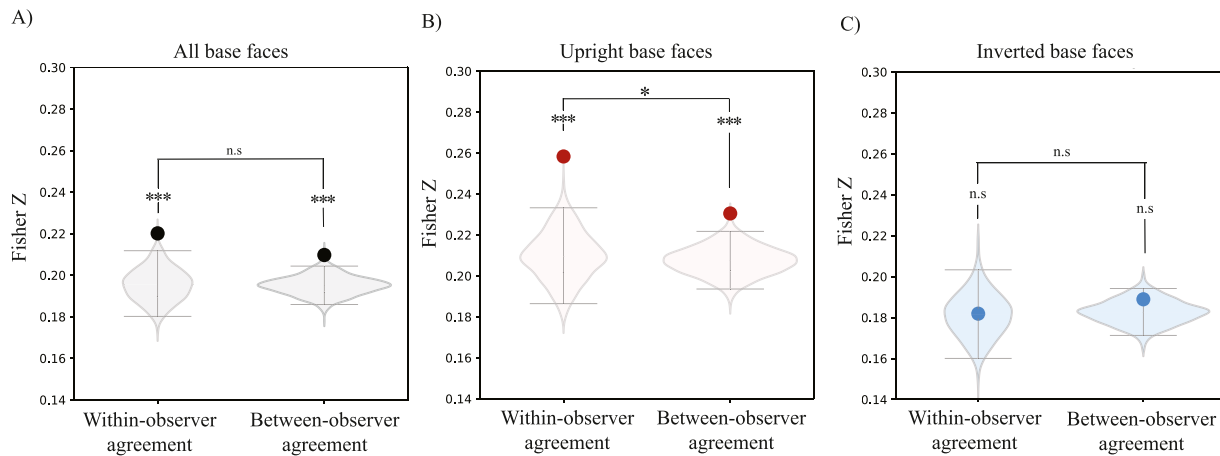


Figure 4. Results of [Experiment 1](#). (A) Within- and between-observer agreement for classification images collapsed across upright and inverted faces. (B) Within- and between-observer agreement for classification images of upright base Mooney faces. (C) Within- and between-observer agreement for classification images of inverted base Mooney faces. In all panels, solid dots represent the empirical correlations, and the violin plots represent the respective permuted null distributions. Error bars represent the upper 97.5% and lower 2.5% boundaries of the permuted null correlation distribution. \*\*\* $p < 0.001$ , \*\* $p < 0.01$ , \* $p < 0.05$ . Nonsignificant comparisons are indicated by n.s.

classification images represented the observer's underlying face representations. To test whether classification images were even nonrandom, we collapsed the upright and inverted faces (vertical pairs of faces in [Figures 3A, 3B](#)) together. On the one hand, the collapsed within observer correlation was the average of the within observer correlations of upright and inverted classification images. This averaged correlation allowed us to investigate whether, across all base faces, classification images were not random. On the other hand, the collapse between observer correlations was computed as the average of the between-observer correlations of upright and inverted classification images, and it represented the agreement across observers for all faces. Altogether, the goal of this analysis was to check and confirm that we could even measure classification images that were significantly different than random responses. If classification images of face-likeness can be measured, then there should be significant within- and between-subject consistency. We found that there was both a significant within- and between-observer correlation for collapsed classification images of Mooney faces (within: Fisher  $Z = 0.22$ ,  $p < 0.01$ ; between: Fisher  $Z = 0.21$ ,  $p < 0.01$ ; [Figure 4A](#)). This result shows that we effectively measured classification images of face-likeness and that they were not due to random responses.

The Mooney faces that we used in this experiment require holistic processing to be recognized ([Canas-Bajo & Whitney, 2020; 2022; Farzin et al., 2009; Latinus & Taylor, 2005; McKone, 2004; Moore & Cavanagh, 1998](#)). If recognition of Mooney faces is supported by holistic face templates, then classification images of Mooney faces should reflect an advantage for

upright faces. Inverted Mooney faces are hard or even impossible to recognize ([Kanwisher, Tong & Nakayama, 1998; McKone, 2004; Latinus and Taylor, 2005](#)), so we expected weaker classification images of inverted Mooney faces. [Figure 3](#) seems to support this division, showing what appear to be weaker classification images for inverted faces. To address this question quantitatively, we computed the within- and between-observer agreement for classification images of upright and inverted faces, separately. We found that there was a significant within-observer agreement in classification images for upright faces (Fisher  $Z = 0.26$ ,  $p < 0.001$ , per permutation test; [Figure 4B](#)) but not inverted faces (Fisher  $Z = 0.18$ ,  $p > 0.05$ , per permutation test; [Figure 4C](#)). A direct comparison of upright and inverted within-observer correlations further confirmed that the day-to-day consistency of the classification images was higher for upright than inverted base faces ( $t = 90.96$ ,  $p < 0.001$ ). That is, only upright Mooney base faces led to non-random, consistent classification images from session to session. Inverted Mooney faces did not. Similarly, we found a significant between-observer agreement only for upright Mooney faces (Fisher  $Z = 0.23$ ,  $p < 0.01$ ; [Figure 4B](#)), but not inverted Mooney faces (Fisher  $Z = 0.19$ ,  $p > 0.05$ ; [Figure 4C](#)).

One of our primary goals was to investigate whether classification images of Mooney faces were observer specific. To this end, we compared the within- and between- observer agreement in the classification images of upright and inverted faces, separately. We found that within-observer agreement of classification images of upright Mooney faces was significantly higher than the between-observer agreement ( $p < 0.05$ ,

per permutation test; Figure 4B). This difference was not significant in classification images of inverted Mooney faces ( $p > 0.05$ , per permutation test; Figure 4C). This result suggests that there are day-to-day consistent individual differences in the classification images of upright Mooney faces, which require holistic processing.

## Discussion

The results of Experiment 1 indicate that classification images of Mooney faces are consistent from session to session. We also found that there are individual differences in classification images of Mooney faces, specific to upright faces. This result is consistent with the idea that holistic face recognition is supported by matching Mooney faces to face templates (Cavanagh, 1991) and that these templates are idiosyncratic—unique to each individual observer. One concern with Experiment 1 is that data were collected online using the Qualtrics platform, and participants completed the experiment on their own computer monitors. Thus we could not gamma correct or control the linearity of each individual participant's monitor. Consequently, the individual differences found may have been influenced by differences across monitors. It is important to note that our inverted condition served as a partial control for this possibility, and the dissociation in results found for upright and inverted faces suggests that the individual differences found in upright Mooney faces are probably not due to differences in gamma correction across the monitors. Nevertheless, individual differences in monitors are an important concern, as are individual differences in room set up, ambient lighting, monitor distance, and other environmental factors. We therefore conducted a second experiment, in the laboratory, with a single controlled and calibrated monitor. The goal of Experiment 2 was to replicate Experiment 1 results and to further investigate whether the individual differences in classification images reflect an underlying idiosyncrasy in the internal representations of faces.

## Experiment 2

Experiment 2 had two goals. First, to control for a possible contribution of individual differences in monitors, hardware, software settings, environmental factors, or any other differences due to the online nature of the first experiment. We therefore conducted the experiment in-person with a controlled environment, constant hardware, and gamma correction of the monitor used for all subjects. The second, more theoretically important goal of Experiment 2 was to investigate whether there is pattern completion within

classification images of Mooney faces. If Mooney faces are recognized through a matching process to stored face templates, then classification images should show that observers fill in missing face feature information within the original Mooney face. Additionally, if holistic face recognition is supported by idiosyncratic templates as suggested by Experiment 1, then we should expect that the way information is filled in is unique to each observer, at least for upright faces.

## Method

### Participants

Nineteen participants (eight male, eleven female) took part in this experiment. All subjects were undergraduate students at the University of California, Berkeley, and provided written consent forms before participation. All experimental procedures were approved by the UC Berkeley Institutional Review Board.

### Material

Stimuli were presented on a gamma corrected CRT monitor at 100 Hz refresh rate, with  $1024 \times 768$  pixels resolution and a horizontal screen size of 40.5cm. The monitor was placed 60 cm from a chin rest that stabilized the participant's head. At this distance, all the face stimuli shown during the experiment subtended  $6^\circ$  visual angle. The presentation of the stimuli was controlled using MATLAB R2016b with Psychophysics Toolbox 3 (Brainard and Vision, 1997; Kleiner, Brainard, & Pelli, 2007).

### Design, procedure and data analysis

Design, procedure, and data analysis was the same as in Experiment 1. Average time between participants' experimental sessions was 1.7 days (Standard Deviation = 0.8)

## Results

The first goal of Experiment 2 was to replicate Experiment 1's results with a single gamma-controlled monitor and consistent environment across observers. Consistent with Experiment 1, we found significant within- and between-observer agreement in classification images collapsed across upright and inverted conditions, which further confirms that that classification images of Mooney faces are consistent from session to session and not due to random noise (within correlation: Fisher  $Z = 0.23$ ,  $p < 0.001$ ; between-subject correlation: Fisher  $Z = 0.22$ ,  $p < 0.001$ ; Figure 5A).



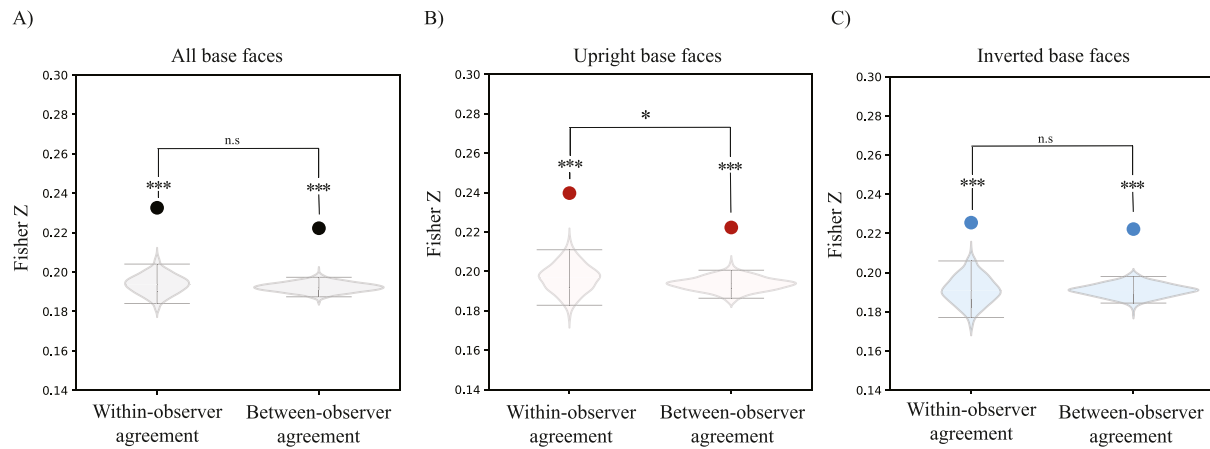


Figure 5. Results of [Experiment 2](#). (A) Within and between-observer agreement for classification images of all upright and inverted base faces collapsed. (B) Within and between-observer agreement for classification images of upright base faces. (C) Within and between-observer agreement for classification images of inverted base faces. In A-B, black dots represent the empirical correlation found and the violin plots represent the respective permuted null correlations. Error bars represent the upper 97.5% and lower 2.5% boundaries of the permuted null correlation distribution. \*\*\* $p < 0.001$ , \*\* $p < 0.01$ , \* $p < 0.05$ . Nonsignificant comparisons are represented with n.s.




Figure 6. Sample resulting classification images of Mooney faces in [Experiment 2](#). (A) Top row: upright Mooney faces (in red). The Mooney face Classification images are recognizable (albeit noisy and ghostly). (B) Bottom row: inverted Mooney faces (in blue). Inverted Mooney face classification images did not appear to be very obvious. See main text for details.

We next compared the consistency and uniqueness of classification images of upright and inverted faces, separately. As in [Experiment 1](#), classification images of upright faces displayed a significant within- and between-observer agreement (within-observer Fisher  $Z = 0.24$ ,  $p < 0.001$ ; between-observer Fisher  $Z = 0.22$ ,  $p < 0.001$ ; [Figure 5B](#)). In this experiment, we found that the within- and between-observer agreement remained significant for inverted faces (within-observer

Fisher  $Z = 0.22$ ,  $p < 0.001$ ; between-observer Fisher  $Z = 0.22$ ,  $p < 0.001$ , per permutation test; [Figure 5C](#)), but it was nevertheless weaker than classification images of upright faces ( $t = 26.32$ ,  $p < 0.001$ ). Indeed, as in [Experiment 1](#), casual inspection also suggests that the classification images of upright faces seem to be more visible and face-like ([Figure 6A](#)) than classification images of inverted faces ([Figure 6B](#)).




Figure 7. (A) Original Mooney face used as base face. (B) In red, original black areas in classification images. Correlations calculated within these signal-free and otherwise featureless regions to analyze the extent of pattern completion. (C) Classification image reflecting pattern completion. (D) Classification image with enhanced contrast to facilitate visualization of the filled-in areas.

Importantly, we found a relatively higher within-observer agreement than between-observer agreement that was specific to upright faces ( $p < 0.05$ , per permutation test; Figure 5B). Replicating Experiment 1, this finding indicates that there are individual differences in classification images that are specific to upright Mooney faces, which require holistic processing. The individual differences found in Experiment 1 are evidently not due to the lack of a consistent environment across observers. In contrast, classification images of inverted faces were not found to show observer-specific idiosyncrasies ( $p > 0.05$ , per permutation test; Figure 5C).

So far, we showed that there are individual differences in classification images of Mooney faces that are consistent session to session. However, the question still remained as to what these day-to-day consistent classification images represent. Our hypothesis was that they reflected participants' internal representation of faces. Mooney faces do not contain easily parsed or bottom-up segmentable information about the parts of the faces (one has to know it is a face before localizing or identifying a particular eye or nose), and some or all face features can be missing from them (e.g., a half-illuminated faces). If classification images reflect participants' templates for faces, then they might reveal filled-in information that is missing at the image level. For example, a base Mooney face that is only partially illuminated on the left side, and mostly in complete darkness on the right side (Figure 7A), might nevertheless evoke a classification image that includes details in the invisible portion of the image. To test this hypothesis, we investigated whether there is pattern completion in classification images of Mooney faces. We operationalized pattern completion as the information present in the classification image only within the black regions of the Mooney face (regions highlighted in red in Figure 7B). Any filling-in or pattern completion of information could involve a number of top-down decision processes and may

be, essentially, an “expectation” that the observer had about whether the face had two eyes, whether it remained obscured in shadow, whether it had an attached contour corresponding to the cheek, and more. To increase the power of the pattern completion analysis and considering the similarity in results of Experiments 1 and 2, we collapsed participant data from both experiments for the following tests.

To measure whether classification images showed pattern completion of missing face features within the black regions of the original Mooney face, we calculated the within-observer agreement within these black regions (Figure 8A). These black regions do not contain any signal, so any consistency in the selected noise patterns within these areas would suggest that the observer has perceptually completed parts of the face that were not originally there. We found significant within- and between-observer agreement within the black regions of all collapsed classification images (within-observer correlation: Fisher  $Z = 0.03$ ,  $p < 0.001$ ; between-observer correlation: Fisher  $Z = 0.02$ ,  $p < 0.001$ ; Figure 8A). This result indicates there is pattern completion in classification images of Mooney faces and is consistent with the hypothesis that recognition of Mooney faces is supported by a matching process to stored templates (Cavanagh, 1991).

When we explored pattern completion within classification images of upright and inverted Mooney faces separately, we found a significant within-subject correlation within the black regions of upright Mooney faces (Fisher  $Z = 0.04$ ,  $p < 0.001$ , per permutation analysis; Figure 8B). This indicates that there was unique filling-in of missing information within single individual observers. Moreover, this pattern completion was consistent from session to session. This confirms that the pattern completion result found in upright Mooney faces was not random; that is, there was meaningful information present in those regions that was not originally there. As a visualization, note a representative observers' classification image in which

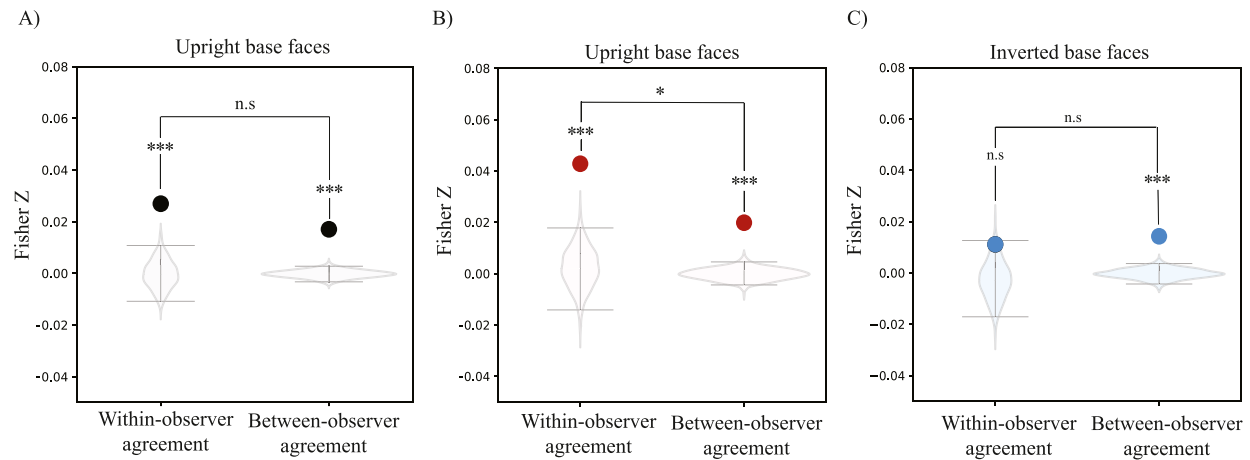


Figure 8. Pattern completion results. (A) Within and between observer agreement in signal-free black areas of all original upright and inverted Mooney faces collapsed. (B) Within and between observer agreement in black areas of the original upright Mooney face. (C) Within and between observer agreement in black areas of the original inverted Mooney face. In all panels, solid dots represent the empirical correlation found and the violin plots represent the respective permuted null correlations. Error bars represent the upper 97.5% and lower 2.5% boundaries of the permuted null correlation distribution. \*\*\* $p < 0.001$ , \*\* $p < 0.01$ , \* $p < 0.05$ . Nonsignificant comparisons are represented with n.s.

pattern completion is evident (Figure 7): even though there is no single face feature information on the right side of the original upright Mooney face (Figure 7A), the classification image of a sample observer showed qualitative completion of the missing features (the right eye, the chin line, the continuation of the mouth, etc.; Figure 7C). Pattern completion becomes more evident visually when we enhance the contrast of the image (Figure 7D). We found no consistent pattern completion for inverted Mooney faces (Fisher  $Z = 0.02$ ,  $p > 0.05$ ; Figure 8C).

Moreover, if observers' templates that support face recognition are idiosyncratic, then classification images should show unique pattern completion for each participant. We found a significant but relatively low between-observer agreement in the way black regions in upright Mooney faces were filled-in (Fisher  $Z = 0.02$ ,  $p < 0.001$ ; Figure 8B). This indicates some consistency between observers, but the within-observer correlation was significantly higher than the between-observer correlation ( $p < 0.5$ , permutation analysis; Figure 8B). These results suggest that there is pattern completion in high-level object representations of upright faces, and that the perceptually filled-in information was unique to each observer. Inverted faces were different. Although we found significant between-observer agreement in the pattern completion of inverted Mooney faces (Fisher  $Z = 0.01$ ,  $p < 0.001$ ; Figure 8C), the difference between within- and between-observer agreement was not significant. This result suggests that there is weaker and less idiosyncratic pattern completion for inverted faces.

## Discussion

In Experiment 2, we aimed to replicate the results found in Experiment 1 under a more controlled experiment set up. In sum, we found that classification images of Mooney faces were consistent from session to session, as in Experiment 1. Similarly, we found observer-specific differences in classification images of Mooney faces that were specific to upright faces. This finding confirms that the results of Experiment 1 were not due to a floor effect.

Furthermore, to understand whether Mooney face recognition is supported by matching them to stored face templates, we explored whether there was day-to-day consistent and unique filling-in of missing facial information in classification images. To address this question, we looked at the within- and between-subject agreement only within the black regions of the original Mooney face. We found within-observer agreement within these black regions, indicating that there is pattern completion in high-level object representations. This result is in line with previous research showing that Mooney face perception is supported by matching to stored templates (Cavanagh, 1991; Damasio, Damasio & Van Hoesen, 1982; Sekuler & Abrams, 1968; Sergent, 1984; Simpson & Crandall, 1972; Smith & Nielsen, 1970; Valentine, 1991). Pattern completion was stronger in upright than inverted classification images. Interestingly, the filled-in pattern was unique to each observer only in upright faces, which suggests that the idiosyncrasy in face recognition is specific to holistic processing.

## General discussion

Previous work suggests that face recognition is supported by a matching process between observed faces and internal face templates (Cavanagh, 1991; Damasio, Damasio & Van Hoesen (1982); Sekuler & Abrams, 1968; Sergent, 1984; Simpson & Crandall, 1972; Smith & Nielsen, 1970; Valentine, 1991). Faces that best match those internal face representations are processed more efficiently than those that do not (Rossion, 2008; Rossion & Boremanse, 2008). However, proposed definitions of face templates often describe a representation of low-level features, which do not generally account for holistic processing of faces and the advantage it affords face recognition (Dakin & Watt, 2009; Viola & Jones, 2004). In the present study, we used classification images of Mooney faces to investigate whether there are face templates that are consistent session to session, whether they support holistic processing of Mooney faces, whether they are observer specific, and whether they allow for pattern completion of holistic representations.

Our results show that there are day-to-day consistent classification images of Mooney faces. The reverse-correlation task used here was purely data driven: observers were not biased toward any external definition of face-likeness, so they applied their internal face presentations to classify the stimuli as face-like. Note that unlike face discrimination and identification tasks commonly used in classification image experiments, the face-likeness task used here may also involve detection processes (Tyler, & Chen, 2006). This suggests that observers' classification images reflected their underlying, implicit face representations. Importantly, we found that classification images of upright Mooney, but not inverted Mooney faces, were observer specific. This is consistent with the idea that there are stored face templates that support holistic processing and that those templates may be idiosyncratic, which could be driving the individual differences in holistic processing found previously in the literature (Canas-Bajo & Whitney, 2020; de Heering & Rossion, 2008; Goodman, Sayfan, Lee, Sandhei, Walle-Olsen, Magnussen, Pezdek, & Arredondo, 2007; Ferguson, Kulkofsky, Cashon & Casasola, 2009; Kuefner, Cassia, Vescovo & Picozzi, 2010; Susilo, Crookes, McKone, & Turner, 2009). These results support the conclusion that holistic and part-based processing are distinct mechanisms of face recognition (McKone, 2004; Moscovitch, Winocur & Behrmann, 1997), and that holistic processing is idiosyncratic, whereas part-based processing is relatively universal (Canas-Bajo & Whitney, 2020).

Furthermore, we found that observers filled in missing information in the original base Mooney face. Humans are highly sensitive to incomplete and

partial visual information, resulting in object and scene recognition that proceeds seemingly unimpaired by occlusions and noise. Many examples of this have been reported, including contour and surface-based filling-in and pattern completion (Gold, Murray, Bennett, Sekuler, 2000; Sekuler, Gold, Murray & Bennett, 2000). Here, we used Mooney faces as stimuli, which do not contain enough low-level information to be uniquely processed in a part-based manner (Canas-Bajo & Whitney, 2020). The pattern completion results found in the present study cannot be explained by known contour integration, surface filling-in, or feature-based pattern completion processes. Interestingly, we found that the way missing information was filled in was idiosyncratic only in upright Mooney faces. Mooney face recognition may therefore be supported by matching the impoverished images to idiosyncratic stored templates. Future research should investigate whether the filled-in or completed regions correspond to edges of the face (so called “attached-contours,” after discounting shadows) or whether the classification image also reveals cast shadows edges (which are uninformative about the structure of the face; Cavanagh 1991).

Last, we also found weaker but day-to-day consistent classification images of inverted Mooney faces with weaker pattern completion and no idiosyncrasies. The stability of these inverted classification images was less efficient than for upright faces, which confirms that our results were not due to a floor effect in the inverted condition. It also agrees with the idea that the part-based pathway to face recognition may be less efficient (Sekuler, Gaspar, Gold & Bennett, 2004). Mooney faces can occasionally maintain some low-level features and be partially supported by a combination of holistic and part-based processing (Canas-Bajo & Whitney, 2020). This suggests that face templates may not only support holistic processing of upright faces but also part-based processing of inverted faces and perhaps object recognition more generally. This is in line with previous work that reported templates for objects (Biederman & Bar, 1999; Rock & DiVita, 1987; Poggio & Edelman, 1990; Logothetis, Pauls, Bülthoff, & Poggio, 1994; Tarr, 1995; Tarr & Bülthoff, 1995).

Our results show that classification images may reveal information about internal stored templates that are unique to each observer. However, it is worth noting that the classification images we measured may have been influenced by nonlinear and decision-related processes that cannot be discerned from our results (Murray, 2011; Neri, 2018). Future research should quantify to what extent an observer's behavior can be predicted from their unique classification image and to what extent a classification image captures aspects of visual processing that significantly influence an observer's trial-by-trial task behavior. Nevertheless, our results showed that there are significant individual



differences in observers' classification images, which suggest that they may reveal unique information about the templates that observers putatively use to recognize Mooney faces.

A long history of computer vision research has aimed to reproduce the template-matching theory of face recognition. Turk and Pentland (1991) defined a computational face recognition system in which faces are encoded as vectors of weights, called eigenfaces, and stored in memory. Eigenfaces are a low-cost representation of faces. When the system encounters a novel face, recognition occurs through a matching procedure to the stored eigenfaces—a type of template. Others have developed Deep Convolutional Neural Networks models of face recognition that highlight the need for internal representations to be general and flexible, so that they can achieve the level of expertise with familiar and unfamiliar faces that we see in humans (Blauch, Behrmann & Plaut, 2021; Orru, Marcialis & Roli, 2020), parallel to the well-known advantage for familiar faces. This and other computer-vision-based face recognition research suggests that models that use template matching can, at least in principle, reproduce the patterns of human face recognition.

Future research can use our approach to investigate the origin and development of face templates at the observer level. The proposed approach of studying holistic perception using reverse correlation classification image method is powerful and could be especially useful in narrowing and identifying the neural mechanisms of holistic face perception. Unfortunately, it is a very time-consuming method, but it can be done within individual observers, as our data suggests. Future work might therefore be able to use the technique in combination with electroencephalography, functional magnetic resonance imaging, and magnetoencephalography. The technique might also be useful for evaluating the biological plausibility and correspondence between face identification algorithms and holistic face perception by assessing classification images in both domains.

Altogether, our results provide evidence that classification images can be used as an unbiased method to tap into observer-specific internal representations of faces. Furthermore, the present study provides a new method for future research to investigate the nature of face templates, such as understanding how many templates are required to support our expert face recognition system (Laurence, Baker, Proietti & Mondloch, 2021); investigating the plasticity of face templates over development (de Heering & Rossion, 2008; Laurence, Baker, Proietti & Mondloch, 2021); or making predictions from an observer's classification image of which faces will be more efficiently processed.

*Keywords: individual differences, holistic perception, face recognition, classification images*

## Acknowledgments

The authors are grateful for helpful discussions with Ervin Hafter, Ken Nakayama, and William Prinzmetal. The authors also thank Maya Zhu, Aidan Gauper and Ethan Shedd for their assistance in data collection.

Supported in part by United States-Spanish Fulbright Commission and NIH grant RO1CA236793.

Commercial relationships: none.

Corresponding author: Teresa Canas-Bajo.

Email: teresa\_canasbajo@berkeley.edu.

Address: Vision Science Graduate Group, University of California, Berkeley, CA, USA.

## References

- Ahumada, A. J., Jr. (1996). Perceptual classification images from Vernier acuity masked by noise. *Perception*, 25(1\_suppl), 2–2.
- Ahumada, A. J. (2002). Classification image weights and internal noise level estimation. *Journal of Vision*, 2(1), 8–8.
- Ahumada, A. Jr., & Lovell, J. (1971). Stimulus features in signal detection. *The Journal of the Acoustical Society of America*, 49(6B), 1751–1756.
- Andrews, T. J., & Schluppeck, D. (2004). Neural responses to Mooney images reveal a modular representation of faces in human visual cortex. *Neuroimage*, 21(1), 91–98.
- Beard, B. L., & Ahumada, A. J., Jr (1998). Technique to extract relevant image features for visual tasks. In *Human Vision and Electronic Imaging III* (Vol. 3299, pp. 79–85). Chicago: SPIE.
- Biederman, I., & Bar, M. (1999). One-shot viewpoint invariance in matching novel objects. *Vision Research*, 39(17), 2885–2899.
- Bindemann, M., Burton, A. M., Hooge, I. T., Jenkins, R., & De Haan, E. H. (2005). Faces retain attention. *Psychonomic Bulletin & Review*, 12(6), 1048–1053.
- Blauch, N. M., Behrmann, M., & Plaut, D. C. (2021). Computational insights into human perceptual expertise for familiar and unfamiliar face recognition. *Cognition*, 208, 104341.
- Bona, S., Cattaneo, Z., & Silvanto, J. (2016). Investigating the causal role of rOFA in holistic detection of mooney faces and objects: an fMRI-guided TMS study. *Brain Stimulation*, 9(4), 594–600.
- Boucart, M., Lenoble, Q., Quetelart, J., Szaffarczyk, S., Desprez, P., & Thorpe, S. J. (2016). Finding

- faces, animals, and vehicles in far peripheral vision. *Journal of Vision*, 16(2), 10–10.
- Brainard, D. H., & Vision, S. (1997). The psychophysics toolbox. *Spatial vision*, 10(4), 433–436.
- Brinkman, L., Dotsch, R., Zondergeld, J., Koevoets, M. G., Aarts, H., & van Haren, N. E. (2019). Visualizing mental representations in schizophrenia patients: A reverse correlation approach. *Schizophrenia Research: Cognition*, 17, 100138.
- Brinkman, L., Todorov, A., & Dotsch, R. (2017). Visualising mental representations: A primer on noise-based reverse correlation in social psychology. *European Review of Social Psychology*, 28(1), 333–361.
- Brooks, J. A., & Freeman, J. B. (2018). Conceptual knowledge predicts the representational structure of facial emotion perception. *Nature Human Behaviour*, 2(8), 581–591.
- Brooks, J. A., Stolier, R. M., & Freeman, J. B. (2018). Stereotypes bias visual prototypes for sex and emotion categories. *Social Cognition*, 36(5), 481–493.
- Brodski, A., Paasch, G. F., Helbling, S., & Wibral, M. (2015). The faces of predictive coding. *Journal of Neuroscience*, 35(24), 8997–9006.
- Bruce, V., & Young, A. (1986). Understanding face recognition. *British Journal of Psychology*, 77(3), 305–327.
- Burton, A. M., Bruce, V., & Hancock, P. J. (1999). From pixels to people: A model of familiar face recognition. *Cognitive Science*, 23(1), 1–31.
- Canas-Bajo, T., & Whitney, D. (2020). Stimulus-specific individual differences in holistic perception of Mooney faces. *Frontiers in Psychology*, 11, 585921.
- Canas-Bajo, T., & Whitney, D. (2022). Relative tuning of holistic face processing towards the fovea. *Vision Research*. 197, 108049.
- Cavanagh, P. (1991). What's up in top-down processing. *Representations of Vision: Trends and Tacit Assumptions in Vision Research*, 295–304.
- Cerf, M., Frady, E. P., & Koch, C. (2009). Faces and text attract gaze independent of the task: Experimental data and computer model. *Journal of Vision*, 9(12), 10–10.
- Cerf, M., Harel, J., Huth, A., Einhäuser, W., & Koch, C. (2008, May). Decoding what people see from where they look: Predicting visual stimuli from scanpaths. In *International Workshop on Attention in Cognitive Systems* (pp. 15–26). Berlin, Heidelberg: Springer.
- Costela, F. M., & Woods, R. L. (2019). When watching video, many saccades are curved and deviate from a velocity profile model. *Frontiers in Neuroscience*, 12, 960.
- Coutrot, A., & Guyader, N. (2014, October). An audiovisual attention model for natural conversation scenes. In *2014 IEEE International Conference on Image Processing (ICIP)* (pp. 1100–1104). IEEE.
- Coutrot, A., & Guyader, N. (2014). How saliency, faces, and sound influence gaze in dynamic social scenes. *Journal of Vision*, 14(8), 5–5.
- Crouzet, S. M., Kirchner, H., & Thorpe, S. J. (2010). Fast saccades toward faces: face detection in just 100 526ms. *Journal of Vision*, 10(4), 16–16.
- Dakin, S. C., & Watt, R. J. (2009). Biological “bar codes” in human faces. *Journal of Vision*, 9(4), 2–2.
- Damasio, A. R., Damasio, H., & Van Hoesen, G. W. (1982). Prosopagnosia: anatomic basis and behavioral mechanisms. *Neurology*, 32(4), 331–331.
- DeGutis, J., Mercado, R. J., Wilmer, J., & Rosenblatt, A. (2013). Individual differences in holistic processing predict the own-race advantage in recognition memory. *PLoS One*, 8(4), e58253.
- DeGutis, J., Wilmer, J., Mercado, R. J., & Cohan, S. (2013). Using regression to measure holistic face processing reveals a strong link with face recognition ability. *Cognition*, 126(1), 87–100.
- De Heering, A., & Rossion, B. (2008). Prolonged visual experience in adulthood modulates holistic face perception. *PloS one*, 3(5), e2317.
- Dotsch, R., & Todorov, A. (2012). Reverse correlating social face perception. *Social Psychological and Personality Science*, 3(5), 562–571.
- Dotsch, R., Wigboldus, D. H., Langner, O., & van Knippenberg, A. (2008). Ethnic out-group faces are biased in the prejudiced mind. *Psychological Science*, 19(10), 978–980.
- Dotsch, R. O. N., Wigboldus, D. H., & Van Knippenberg, A. D. (2013). Behavioral information biases the expected facial appearance of members of novel groups. *European Journal of Social Psychology*, 43(1), 116–125.
- Éthier-Majcher, C., Joubert, S., & Gosselin, F. (2013). Reverse correlating trustworthy faces in young and older adults. *Frontiers in Psychology*, 4, 592.
- Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. N. (1998). What is “special” about face perception? *Psychological Review*, 105(3), 482.
- Farzin, F., Rivera, S. M., & Whitney, D. (2009). Holistic crowding of Mooney faces. *Journal of Vision*, 9(6), 18–18.
- Ferguson, K. T., Kulkofsky, S., Cashon, C. H., & Casasola, M. (2009). The development of specialized processing of own-race faces in infancy. *Infancy*, 14(3), 263–284.

- Foulsham, T., Cheng, J. T., Tracy, J. L., Henrich, J., & Kingstone, A. (2010). Gaze allocation in a dynamic situation: Effects of social status and speaking. *Cognition*, *117*(3), 319–331.
- Frank, M. C., Vul, E., & Johnson, S. P. (2009). Development of infants' attention to faces during the first year. *Cognition*, *110*(2), 160–170.
- Gauthier, I. (2020). What we could learn about holistic face processing only from nonface objects. *Current Directions in Psychological Science*, *29*(4), 419–425.
- Gold, J. M., Murray, R. F., Bennett, P. J., & Sekuler, A. B. (2000). Deriving behavioural receptive fields for visually completed contours. *Current Biology*, *10*(11), 663–666.
- Goodman, G. S., Sayfan, L., Lee, J. S., Sandhei, M., Walle-Olsen, A., & Magnussen, S., ... & Arredondo, P. (2007). The development of memory for own-and other-race faces. *Journal of Experimental Child Psychology*, *98*(4), 233–242.
- Goold, J. E., & Meng, M. (2016). Visual search of Mooney faces. *Frontiers in Psychology*, *7*, 155.
- Gosselin, F., & Schyns, P. G. (2003). Superstitious perceptions reveal properties of internal representations. *Psychological Science*, *14*(5), 505–509.
- Hole, G. J., George, P. A., Eaves, K., & Rasek, A. (2002). Effects of geometric distortions on face-recognition performance. *Perception*, *31*(10), 1221–1240.
- Jack, R. E., Caldara, R., & Schyns, P. G. (2012). Internal representations reveal cultural diversity in expectations of facial expressions of emotion. *Journal of Experimental Psychology: General*, *141*(1), 19.
- Jack, R. E., & Schyns, P. G. (2015). The human face as a dynamic tool for social communication. *Current Biology*, *25*(14), R621–R634.
- Kanwisher, N., Tong, F., & Nakayama, K. (1998). The effect of face inversion on the human fusiform face area. *Cognition*, *68*(1), B1–B11.
- Karremans, J. C., Dotsch, R., & Corneille, O. (2011). Romantic relationship status biases memory of faces of attractive opposite-sex others: evidence from a reverse-correlation paradigm. *Cognition*, *121*(3), 422–426.
- Ke, T. W., Stella, X. Y., & Whitney, D. (2017, September). Mooney face classification and prediction by learning across tone. In *2017 IEEE International Conference on Image Processing (ICIP)* (pp. 2025–2029). IEEE.
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3? *Perception*, *36*, 4.
- Kuefner, D., Cassia, V. M., Vescovo, E., & Picozzi, M. (2010). Natural experience acquired in adulthood enhances holistic processing of other-age faces. *Visual Cognition*, *18*(1), 11–25.
- Latinus, M., & Taylor, M. J. (2005). Holistic processing of faces: learning effects with Mooney faces. *Journal of Cognitive Neuroscience*, *17*(8), 1316–1327.
- Laurence, S., Baker, K. A., Proietti, V. M., & Mondloch, C. J. (2021). What happens to our representation of identity as familiar faces age? Evidence from priming and identity aftereffects. *British Journal of Psychology*, *113*(3), 677–695.
- Logothetis, N. K., Pauls, J., Bülthoff, H. H., & Poggio, T. (1994). View-dependent object recognition by monkeys. *Current Biology*, *4*(5), 401–414.
- Mangini, M. C., & Biederman, I. (2004). Making the ineffable explicit: Estimating the information employed for face classifications. *Cognitive Science*, *28*(2), 209–226.
- Marat, S., Rahman, A., Pellerin, D., Guyader, N., & Houzet, D. (2013). Improving visual saliency by adding “face feature map” and “center bias.” *Cognitive Computation*, *5*(1), 63–75.
- Mareschal, I., Dakin, S. C., & Bex, P. J. (2006). Dynamic properties of orientation discrimination assessed by using classification images. *Proceedings of the National Academy of Sciences*, *103*(13), 5131–5136.
- Martin, J. G., Davis, C. E., Riesenhuber, M., & Thorpe, S. J. (2018). Zapping 500 faces in less than 100 seconds: evidence for extremely fast and sustained continuous visual search. *Scientific Reports*, *8*(1), 1–12.
- Maurer, D., Le Grand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, *6*(6), 255–260.
- McKone, E. (2004). Isolating the special component of face recognition: peripheral identification and a Mooney face. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(1), 181.
- Moore, C., & Cavanagh, P. (1998). Recovery of 3D volume from 2-tone images of novel objects. *Cognition*, *67*(1-2), 45–71.
- Mooney, C. M. (1957). Age in the development of closure ability in children. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, *11*(4), 219.
- Moscovitch, M., Winocur, G., & Behrmann, M. (1997). What is special about face recognition? Nineteen experiments on a person with visual object agnosia and dyslexia but normal face recognition. *Journal of Cognitive Neuroscience*, *9*(5), 555–604.
- Murray, R. F. (2011). Classification images: A review. *Journal of Vision*, *11*(5), 2–2.

- Neri, P. (2018). Classification images as descriptive statistics. *Journal of Mathematical Psychology*, 82, 26–37.
- Neri, P., & Heeger, D. J. (2002). Spatiotemporal mechanisms for detecting and identifying image features in human vision. *Nature Neuroscience*, 5(8), 812–816.
- Orrù, G., Marcialis, G. L., & Roli, F. (2020). A novel classification-selection approach for the self-updating of template-based face recognition systems. *Pattern Recognition*, 100, 107121.
- Pascalis, O., & Kelly, D. J. (2009). The origins of face processing in humans: Phylogeny and ontogeny. *Perspectives on Psychological Science*, 4(2), 200–209.
- Poggio, T., & Edelman, S. (1990). A network that learns to recognize three-dimensional objects. *Nature*, 343(6255), 263–266.
- Ratner, K. G., Dotsch, R., Wigboldus, D. H., van Knippenberg, A., & Amodio, D. M. (2014). Visualizing minimal ingroup and outgroup faces: Implications for impressions, attitudes, and behavior. *Journal of Personality and Social Psychology*, 106(6), 897.
- Ringach, D., & Shapley, R. (2004). Reverse correlation in neurophysiology. *Cognitive Science*, 28(2), 147–166.
- Ro, T., Russell, C., & Lavie, N. (2001). Changing faces: A detection advantage in the flicker paradigm. *Psychological Science*, 12(1), 94–99.
- Rock, I., & DiVita, J. (1987). A case of viewer-centered object perception. *Cognitive Psychology*, 19(2), 280–293.
- Rossion, B. (2008). Picture-plane inversion leads to qualitative changes of face perception. *Acta Psychologica*, 128(2), 274–289.
- Rossion, B., & Boremanse, A. (2008). Nonlinear relationship between holistic processing of individual faces and picture-plane rotation: Evidence from the face composite illusion. *Journal of Vision*, 8(4), 3–3.
- Russell, R., Duchaine, B., & Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychonomic Bulletin & Review*, 16(2), 252–257.
- Schwiedrzik, C. M., Melloni, L., & Schurger, A. (2018). Mooney face stimuli for visual perception research. *PLoS One* 13, e0200106, doi:10.1371/journal.pone.0200106.
- Sekuler, R. W., & Abrams, M. (1968). Visual sameness: a choice time analysis of pattern recognition processes. *Journal of Experimental Psychology*, 77(2), 232.
- Sergent, J. (1984). An investigation into component and configural processes underlying face perception. *British Journal of Psychology*, 75(2), 221–242.
- Sekuler, A. B., Gaspar, C. M., Gold, J. M., & Bennett, P. J. (2004). Inversion leads to quantitative, not qualitative, changes in face processing. *Current Biology*, 14(5), 391–396.
- Sekuler, A. B., Gold, J. M., Murray, R. F., & Bennett, P. (2000). Visual completion of partly occluded objects: Insights from behavioral studies. *Neuroophthalmology*, 23, 165–168.
- Simpson, W. E., & Crandall, S. J. (1972). The perception of smiles. *Psychonomic Science*, 29(4), 197–200.
- Smith, E. E., & Nielsen, G. D. (1970). Representations and retrieval processes in short-term memory: Recognition and recall of faces. *Journal of Experimental Psychology*, 85(3), 397.
- Solomon, J. A. (2002). Noise reveals visual mechanisms of detection and discrimination. *Journal of Vision*, 2(1), 7–7.
- Susilo, T., Crookes, K., McKone, E., & Turner, H. (2009). The composite task reveals stronger holistic processing in children than adults for child faces. *PLoS One*, 4(7), e6460.
- Tarr, M. J. (1995). Rotating objects to recognize them: A case study on the role of viewpoint dependency in the recognition of three-dimensional objects. *Psychonomic Bulletin & Review*, 2(1), 55–82.
- Tarr, M. J., & Bülthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or by multiple views? Comment on Biederman and Gerhardstein. *Journal of Experimental Psychology: Human Perception and Performance*, 21(6), 1494–1505.
- Theeuwes, J., & Van der Stigchel, S. (2006). Faces capture attention: Evidence from inhibition of return. *Visual Cognition*, 13(6), 657–665.
- Todorov, A., Dotsch, R., Wigboldus, D. H., & Said, C. P. (2011). Data-driven methods for modeling social perception. *Social and Personality Psychology Compass*, 5(10), 775–791.
- Turk, M. A., & Pentland, A. P. (1991, January). Face recognition using eigenfaces. In *Proceedings. 1991 IEEE computer society conference on computer vision and pattern recognition* (pp. 586–587). IEEE Computer Society.
- Tyler, C. W., & Chen, C. C. (2006). Spatial summation of face information. *Journal of Vision*, 6(10), 11–11.
- Umiltà, C., Simion, F., & Valenza, E. (1996). Newborn's preference for faces. *European Psychologist*, 1(3), 200–205.



- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *The Quarterly Journal of Experimental Psychology Section A*, 43(2), 161–204.
- Victor, J. D. (2005). Analyzing receptive fields, classification images and functional images: challenges with opportunities for synergy. *Nature Neuroscience*, 8(12), 1651–1656.
- Viola, P., & Jones, M. J. (2004). Robust real-time face detection. *International Journal of Computer Vision*, 57(2), 137–154.
- Wang, R., Li, J., Fang, H., Tian, M., & Liu, J. (2012). Individual differences in holistic processing predict face recognition ability. *Psychological Science*, 23(2), 169–177.
- Watson, A. B. (1998). Multi-category classification: Template models and classification images. *Investigative Ophthalmology and Visual Science*, 39(ARVO Suppl. 4), S912.
- Wolfe, B. A., & Whitney, D. (2015). Saccadic remapping of object-selective information. *Attention, Perception, & Psychophysics*, 77(7), 2260–2269.
- Xia, Y., Manassi, M., Nakayama, K., Zipser, K., & Whitney, D. (2020). Visual crowding in driving. *Journal of Vision*, 20(6), 1–1.
- Young, A. W., & Bruce, V. (2011). Understanding person perception. *British Journal of Psychology*, 102(4), 959–974.