The Face Inversion Effect: Roles of First- and Second-Order Configural Information

Ciro CIVILE, Rossy MCLAREN, and Ian P. L. MCLAREN

University of Exeter

Abstract

The face inversion effect (FIE) is a reduction in recognition performance for inverted faces
compared with upright faces. Several studies have proposed that a type of configural
information, called second-order relational information, becomes more important with increasing
expertise and gives rise to the FIE. However, recently it has been demonstrated that it is possible
to obtain an FIE with facial features presented in isolation, showing that configural information
is not necessary for this effect to occur. In this article we test whether there is a role for
configural information in producing the FIE and whether second- or first-order relational
information is particularly important. In Experiment 1, we investigated the role of configural
information and local feature orientation by using a new type of "Thatcherizing" transformation
on our set of faces, aiming to disrupt second-order and local feature orientation information but
keeping all first-order properties unaltered. The results showed a significant reduction in the FIE
for these "new" Thatcherized faces, but it did not entirely disappear. Experiment 2 confirmed the
FIE for new Thatcherized faces, and Experiment 3 establishes that both local feature orientation
and first-order relational information have a role in determining the FIE.

Key words: inversion effect, configural information, featural information, Thatcher illusion, face
recognition

Face recognition is one of the best cognitive skills people have. We generally recognize faces with very little effort, despite large variations in skin tone, viewpoint, and expression. Discussion of the nature of face perception has been divided into two interpretations: One asserts that a large body of research supports the notion of specialized mechanisms used to process facial stimuli (Rhodes & Tremewan, 1994; Valentine, 1988; Yin, 1969), whereas the other points out that face recognition is actually based on general mechanisms that can also operate for other nonfacial stimuli as well (Diamond & Carey, 1986; Tanaka & Farah, 1991). One of the most robust phenomena, and central to this debate, is the face inversion effect (FIE), which is the disproportionate (relative to other stimuli such as pictures of houses) decrease in recognition performance for upside down (inverted) faces relative to upright faces (Yin, 1969). On its discovery, the FIE was described as a clear consequence of the specialized mechanisms used in face processing, which explained why the impairment in recognizing upside down faces was significantly larger than that for other objects (Yin, 1969).

However, Diamond and Carey (1986) provided a new, alternative account of the FIE suggesting that the inversion effect does not reflect a face-specific process. The authors demonstrated that the inversion effect on recognition memory can be as strong with images of dogs as with faces if the subjects are experts in the identification and assessment of specific dog breeds. This analysis suggests that the only stimuli that result in a substantial inversion effect are ones for which the subjects have the necessary expertise. The authors distinguished between three types of information that can be used in recognition: isolated features and first-order and second-order relational information. Isolated or local features are the independent constituent elements of an object (e.g., the eyes, nose, mouth). First-order information consists of spatial relations between constituent elements of an object (e.g., the arrangement of the nose above the

mouth). It is the first-order information that specifies a set of facial features as a face. Second-order information defines the relative size of these spatial relationships with regard to a base prototype. All faces tend to have the same first-order relational information in common; the essential relational information by which faces differ from each other is second-order. These two kinds of relational information, first- and second-order, are both types of configural information. Diamond and Carey suggested that a large inversion effect will be obtained if just three conditions are met. First, the members of the class of stimuli must share a configuration. Second, it must be possible to identify the individual members of the class through second-order information. Finally, percipients must have the expertise to exploit such second-order information.

This sensitivity of configural information to inversion is also often suggested as the basis (at least in part) of the Thatcher illusion (Thompson, 1980). Here a face has its eyes and mouth rotated 180° in situ relative to the rest of the face. When inverted, the transformed face does not look that unusual, but when upright, it can give the face a quite striking appearance. Here, the illusion seems to depend on the inversion of mouth and eyes within the face being hard to detect when the whole face is inverted. The explanation typically offered is that inversion reduces the use of configural information in the face and promotes a more componential analysis of the features present. In isolation, the mouth and eyes do not look odd, and so they cause no great reaction in the viewer. When the face is shown in its normal orientation, however, we revert to configural processing, and this makes the distortions present in the mouth and eyes stand out, resulting in a strong reaction to the face on the part of most percipients. These results, and others like them, all provide evidence for the powerful effect that relational information has in the processing of upright faces relative to inverted faces (Bartlett & Searcy, 1993; Carey, 1992;

3

Carey, Diamond, & Woods, 1980; Leder & Bruce, 1998, 2000; Searcy & Bartlett, 1996).

However, these results do not directly assess the specific role that first-order and second-order

relational information play in producing the FIE.

Tanaka and Farah (1991) conducted a direct test of the claim that second-order relations

are particularly affected by inversion. The subjects were trained to identify dot patterns that

either shared or did not share a configuration. In the study phase, the subjects were trained to

identify these patterns by male and female names. In the test phase, the subjects were asked to

identify the same patterns seen in the study phase presented in their upright and inverted

orientations. The idea was that patterns that did not share a spatial configuration were

discriminable based on first-order relational information, whereas those that did share a spatial

configuration must be discriminated on the basis of second-order information. The results

indicated that inversion had a similar effect on the two kinds of dot pattern stimuli, which

suggested that second-order information was no more vulnerable to inversion than first-order

information (Tanaka & Farah, 1991). Both types of information, then, might contribute to the

FIE.

In addition, a variety of evidence suggests that featural information may have an

important role in face perception. McKone and Yovel (2009) conducted a meta-analysis that

showed that the inversion effect was not entirely due to changes in the use of configural

information caused by inversion but could also depend on the orientation of individual features.

They argued that local feature information can make a contribution to the FIE that is the equal of

that caused by configural information. They evaluated the claim that perception based on local

feature information shows no or weak inversion effects and found that the evidence does not

support this claim. Their position is strongly supported by Rakover and Teucher (1997), who

found that it is possible to obtain an inversion effect even with facial features presented in isolation, suggesting that configural information is not necessary to obtain such an effect.

Overall these studies suggest that the FIE has contributions from both local features and configural information, but there is uncertainty about the nature of the configural contribution, in particular whether it is caused by first- or second-order information. Consistent with this position, Civile, McLaren, and McLaren (2014) showed that the FIE was completely eliminated only when local feature orientation information, in addition to configural information, was disrupted. They used a new type of transformation similar to Thatcherizing sets of scrambled faces (which they called 50% feature-inverted and scrambled). Faces of this type are reproduced later in this article, but in essence the local features of the face (eyes, nose, mouth, ears) were randomly repositioned (scrambling), and half of them (one eye, one ear, one of nose or mouth) were inverted (50% feature inversion). There was no FIE with these stimuli, but there was a reliable FIE with scrambled faces that had not had any local feature inversion. This established that local feature orientation was sufficient to produce a FIE in the absence of any configural information of the type usually found in faces. These stimuli also had the advantage of addressing a problem with Rakover and Teucher's (1997) procedure, in that they presented a single feature to participants in their study (note that they considered both eyes to be a single feature), which might have allowed participants to imagine it as belonging to a normal face. Consequently, it might have been the memory for this imagined face that led to the inversion effect they reported. Civile, McLaren, et al. (2014) demonstrated that the orientation of isolated features controlled at least part of the FIE even when these features were presented in the context of other facial features.

5

This brings us to the question of whether it is possible to establish the nature of the role configural information plays in generating the inversion effect for faces. There is a great deal of evidence consistent with the proposition that it does play a role, even though the results obtained by Rakover and Teucher (1997) and Civile, McLaren, et al. (2014) suggested that it is not necessary for it to be available to participants to generate the FIE. It would be unwise to jump to the conclusion that configural information is not itself effective in producing the inversion effect, however, without first manipulating configural information while controlling for local feature orientation. Thus, in the experiments we report in this article, we sought to investigate the contribution that first-order and second-order configural information may have in determining the FIE. In Experiment 1, we investigated the role of first- and second-order information by transforming our set of faces in such a way as to control for any effect of local feature orientation on the FIE.

EXPERIMENT 1

METHOD

Materials

In Experiment 1 we used 128 images of male faces (male faces were used because they allowed the inclusion of individual ears in the manipulation as well). The faces were standardized in gray-scale format and cropped around the hairline in Adobe Photoshop. Gimp 2.6 was then used to manipulate the facial features of the 128 stimuli. Examples of the stimuli used are given in Figure 1. The experiment was run with Superlab Version 4.0.7b installed on an iMac computer.

<Insert Figure 1 here about here>

We produced a set of normal faces and a set of what we called "new Thatcherized" faces. The original manipulation, by Thompson (1980), involved rotating the mouth and each of the eyes (individually) by 180°. In our manipulation, we rotated (by 180°) one eye (including eyebrow), one ear, and either the nose or mouth of sets of normal faces. The features that had been rotated were counterbalanced so that we created four different sets or categories of new Thatcherized faces (Categories A, B, C, and D), each represented by a prototype. Exemplar faces drawn from a particular category shared the same orientation of the features with that category's prototype. This manipulation approximately balances the number of features that are upright in a face (whether the face itself is inverted or not), thus controlling for the effect of individual features on inversion to an extent that has been shown to be effective in previous experiments (Civile, McLaren, et al., 2014). The second-order information in the new Thatcherized faces will also be somewhat disrupted as a consequence of this transformation, because we have changed the orientation of half the local features, and this must itself alter the small variations in spatial relationship between each feature. However, the first-order configural information would not have been greatly affected because the rotations were carried out in situ. We also cropped the neck of each face, and both normal and new Thatcherized faces were smoothed to control for effects of luminosity and local information (e.g., skin tone, blemishes). Each face stimulus was presented in four different conditions, counterbalanced across the eight participant groups (i.e., normal upright, normal inverted, new Thatcherized upright, and new Thatcherized inverted). Examples of the stimuli used are given in Figure 1.

Participants

The participants were 24 psychology undergraduates at the University of Exeter. We counterbalanced the study by splitting the participants into eight groups. Each participant group was shown the same 128 faces, but each group saw each face in a different condition.

Procedure

The study consisted of a study phase and an old/new recognition phase. After receiving instructions about how to proceed with the experiment, the participants looked at 64 different faces (presented one at a time in random order) during the study phase. After further instructions, participants were then asked to look at 128 faces (including the 64 seen in the study phase) again presented in a random order. During this old/new recognition phase the participants indicated whether they had seen each face during the study phase. In the study phase each participant was shown four different types of face (two selected from only one category of new Thatcherized faces, as upright and inverted exemplars, and normal faces, also as upright and inverted exemplars) with 16 photos for each face type (giving a total of 64 faces). In the test phase another 64 novel faces from the same four face types were added to this set. Each facial stimulus never appeared in more than one condition at a time during the experiment but served at some point in all conditions.

Trial Structure

The first event participants saw after the instructions consisted of a warning cue (a fixation cross in the center of the screen) presented for 1 s. This was followed by a face, presented for 3 s, then the fixation cross was repeated and another face presented until all 64 facial stimuli had been seen. Once all 64 faces had been shown, the program moved to the next set of instructions, which explained the requirements of the old/new recognition task. Participants were told that they were about to see more faces, presented one at a time in random

order. They were asked to press the "." key if they recognized the face or to press "x" if they did not. Each participant within each participant group was then shown (in a random order) the 64 faces they had already seen, intermixed with another 64 unseen faces. These unseen faces were those from the sets of facial stimuli not used during the study phase.

During the old/new recognition task, after the 1-s warning cue, each facial stimulus was shown for 4 s, and the participant had to respond during this period. If the participant pressed the wrong key (i.e., a key other than "x" or ".") the feedback "Wrong key" was shown for 2 s before the next face appearing on the screen. If the participant was too slow in responding (i.e., took longer than 4 s), the message "Too slow" appeared on the screen. Otherwise. no feedback was given. Because there were 128 faces to observe in the old/new recognition task, three participant breaks were incorporated. These allowed participants to rest their eyes after they had viewed 32 facial stimuli. At the end of the experiment participants were shown a message thanking them for participating. They were then paid and debriefed.

RESULTS

As in the other experiments reported in this article, analysis of latencies does not add to the analysis of the accuracy scores. The mean latencies for each stimulus condition used in this experiment were normal upright, 2.54 s; normal inverted, 2.61 s; new Thatcherized upright, 2.79 s; and new Thatcherized inverted, 2.80 s. The data from all 24 participants were used in the signal detection d' analysis of the recognition task (seen and not seen faces for each face type), where a $d' =$ of 0.00 indicates chance-level performance. We also offer an estimate of effect size based on Cohen's d for all comparisons that are $p < .1$ or better, and we also give the appropriate confidence interval.

ANOVA revealed there was no significant main effect of face type, $F(1, 23) = 1.062$, $\underline{MSE} = 0.437$, $\underline{p} = .313$, but there was a significant main effect of orientation, $F(1, 23) = 20.580$, $\underline{MSE} = 0.222$, $\underline{p} < .001$, $\underline{d} = 1.09$, 95% CI [0.93, 1.25], and a significant interaction between face type and orientation, $F(1, 23) = 7.280$, $\underline{MSE} = 0.044$, $\underline{p} = .013$, $\underline{d} = 0.81$, 95% CI [0.53, 1.09] (Figure 2). Thus, simple effect analyses were conducted showing that there was a strong inversion effect for normal faces, $\underline{t}(23) = 4.232$, $\underline{SE} = 0.170$, $\underline{p} < .001$, $\underline{d} = 1.26$, 95% CI [1.03, 1.49], but no significant inversion effect for new Thatcherized faces, $\underline{t}(23) = 1.424$, $\underline{SE} = 0.107$, $\underline{p} = .16$. Performance in recognizing normal upright faces was significantly better than recognition for upright new Thatcherized ones, $\underline{t}(23) = 2.866$, $\underline{SE} = 0.147$, $\underline{p} = .008$, $\underline{d} = 0.87$, 95% CI [0.68, 1.07] (see Figure 2), but the comparison for inverted faces did not approach significance ($\underline{p} = .45$). Additional analyses were done to test performance against chance for each condition. This analysis was not significant for the inverted normal faces, $\underline{t}(31) = 1.229$, $\underline{SE} = 0.128$, $\underline{p} = .23$, but was in all other cases, with the smallest $\underline{t}$ being for the new Thatcherized inverted faces, $\underline{t}(31) = 2.148$, $\underline{SE} = 0.140$, $\underline{p} = .042$, $\underline{d} = 0.87$, 95% CI [0.68, 1.07].

<Insert Figure 2 about here>

DISCUSSION

Experiment 1 shows a strong FIE for the set of normal (but smoothed) faces and significantly better performance in recognizing upright normal faces compared with new Thatcherized faces in an upright orientation. The main finding in Experiment 1 is that the FIE for sets of new Thatcherized faces is significantly different from that for normal faces and is not itself significant. However, it may be that it has not entirely disappeared, an issue we shall return to in a moment. One possible interpretation of the reduction in the inversion effect for the new Thatcherized faces is that the disruption of the second-order relational information caused by our

transformation of the stimuli is responsible for this effect. But the results from Civile, McLaren, et al.'s (2014) studies suggest that the orientation of individual features is, if anything, more likely to be the factor responsible for bringing about these results. If we consider the performance to both sets of inverted faces, we can see numerically (though not significantly) better performance in recognizing new Thatcherized inverted faces compared with normal inverted faces. This result could be explained by the fact that normal inverted faces have all the local feature orientation information disrupted (i.e., all six main features are presented upside down), whereas new Thatcherized inverted faces still benefit from 50% of the features being presented the right way up. The significant difference in performance on upright faces from the two sets could be explained in a similar fashion as being due to upright normal faces having 100% of their features in an upright orientation, and the two effects in combination would then explain the significant stimulus type **<x>** orientation interaction.

The one difficulty with this analysis, which relies only on the orientation of isolated features, is that there is still a numerical inversion effect for the new Thatcherized faces, one that, if we allowed a one-tailed test, would be marginally significant. There is also something of a tension between these results and those obtained in Civile, McLaren, et al.'s (2014) studies using scrambled faces. In those experiments the configural information was completely disrupted (the features were randomly rearranged within the face), but the inversion effect was not much affected. In this experiment, a manipulation that distorts the stimulus much less severely (at least to the naked eye) has significantly disrupted the inversion effect. In the scrambled faces, if three out of six features were inverted the FIE disappeared. Here, the same manipulation may not have completely eliminated this effect. Thus, in Experiment 2 we made a direct comparison between these two versions of the inversion effect.

EXPERIMENT 2

The purpose of this experiment was to compare the FIE obtained solely with the single feature orientation information (Civile, McLaren, et al., 2014) with that obtained with the new Thatcherized faces (from Experiment 1). It also provided us with an opportunity to test the hypothesis that it is possible to obtain a significant FIE with the new Thatcherized faces.

METHOD

Materials

In Experiment 2 we used the sets of scrambled faces previously used in Civile, McLaren, et al.'s (2014) studies, together with the sets of new Thatcherized faces used in Experiment 1. Each of the four categories of scrambled faces was represented by a particular configuration. We scrambled faces by selecting one feature (e.g., an eye, mouth, nose, or ear) at random, then moving it to the forehead (chosen because this is the widest space inside the face and so can accommodate any feature). After this, a second feature was selected and moved to the space left empty by the first feature, and so on until all six facial features had been moved, but their orientation remained the same (i.e., upright). Within an individual category, all the scrambled faces shared the arrangement of the features in common with the prototype. For example, each face drawn from Category A had the features in the locations shown in Figure 3. The subjects in our experiment were presented with stimuli drawn from only one category of scrambled faces (in upright and inverted orientations) and one category of new Thatcherized faces (also in upright and inverted orientations). The four categories of scrambled and new Thatcherized faces were counterbalanced across the eight participant groups.

<Insert Figure 3 about here>

Participants

A total of 32 psychology undergraduates at the University of Exeter took part in the experiment. We counterbalanced the study, as in Experiment 1, by splitting the participants into eight groups.

Procedure

The experimental procedure used was exactly the same as before. The four sets of scrambled faces and four sets of new Thatcherized faces were counterbalanced across the eight participant groups.

RESULTS

The mean latencies for each stimulus condition were scrambled upright, 2.47 s; scrambled inverted, 2.59 s; new Thatcherized upright, 2.54 s; and new Thatcherized inverted, 2.61 s. The data from all 32 participants were used in the signal detection $\underline{d}'$ analysis. ANOVA revealed a significant main effect of face type, $\underline{F}(1, 31) = 7.618$, $\underline{MSE} = 0.260$, $\underline{p} = .012$, $\underline{d} = 0.61$, 95% CI [0.48, 0.75]. Also, there was a significant main effect of orientation, $\underline{F}(1, 31) = 9.563$, $\underline{MSE} = 0.271$, $\underline{p} = .004$, $\underline{d} = 0.72$, 95% CI [0.58, 0.85]; however, there was no significant interaction between face type and orientation, $\underline{F}(1, 31) = 0.104$, $\underline{MSE} = 0.030$, $\underline{p} = .749$, indicating that the inversion effect was similar in size for both stimulus types. Planned comparisons were conducted, showing a strong inversion effect for scrambled faces, $\underline{t}(31) = 2.562$, $\underline{SE} = 0.122$, $\underline{p} = .015$, $\underline{d} = 0.62$, 95% CI [0.45, 0.80], and a trend that approached significance for new Thatcherized faces, $\underline{t}(31) = 1.953$, $\underline{SE} = 0.132$, $\underline{p} = .06$, $\underline{d} = 0.45$, 95% CI [0.26, 0.65]. Performance in recognizing new Thatcherized upright faces was not significantly better (despite the numerical difference) than recognition for upright scrambled ones, $\underline{t}(31) = 1.644$, $\underline{SE} = 0.129$, $\underline{p} = .11$; however, there was a significant difference in the recognition of new Thatcherized inverted faces (which were better) compared with inverted scrambled faces, $\underline{t}(31) =$

2.225, $\underline{SE}$ = 0.121, $\underline{p}$ = .033, $\underline{d}$ = 0.51, 95% CI [0.33, 0.69] (Figure 4). Additional analyses were done to test performance against chance for each condition. The analysis was not significant in the case of inverted scrambled faces, $\underline{t}$(31) = 0.232, $\underline{SE}$ = 0.092, $\underline{p}$ = .817, but was for the other three conditions, with the smallest $\underline{t}$ for new Thatcherized inverted faces, $\underline{t}$(31) = 3.186, $\underline{SE}$ = 0.091, $\underline{p}$ = .003, $\underline{d}$ = 0.79, 95% CI [0.67, 0.92].

<Insert Figure 4 about here>

DISCUSSION

The results from Experiment 2 confirmed that the full disruption of configural information, while leaving the local feature orientations unaltered, leads to a strongly significant inversion effect for sets of scrambled faces (Civile, McLaren, et al., 2014). Also, once again we find a numerical inversion effect, this time a nearly significant result (one that would have been significant on a one-tailed test), for the new Thatcherized faces. Because we have evidence in this experiment that the inversion effect for new Thatcherized faces is not different from that in our scrambled faces, we were able to perform a Bayesian analysis of both Experiments 1 and 2 using the procedures outlined by Dienes (2011). This used the effect in the scrambled faces as the prior, setting the standard deviation of $\underline{p}$(population value|theory) to the mean for the inversion effect in this condition. We then calculated Bayes factors using a one-tailed distribution for our theory (because it does not predict a reversal of the inversion effect, only that it might be diminished) and a mean of 0. This gave a Bayes factor ($\underline{B}$) of 1.48 for Experiment 1 and 3.78 for Experiment 2. To combine these results we simply multiply the Bayes factors to get an overall $\underline{B}$ of 5.59. This is greater than 3, and so the combined results of our two experiments can be taken as substantial evidence for the claim that there is an inversion effect in the new Thatcherized faces. We will return to this point in the Discussion.

A last comment concerns the overall performance on the two sets of stimuli. There is a hint that preserved first-order relations confer a greater benefit in recognizing the face compared with that due simply to single feature orientation information in the upright faces, but this difference was not significant in this experiment. But this numerical difference is noteworthy in the context of the results for the inverted stimuli, which confirmed the influence of the amount of local feature orientation information. If six features are inverted (as in the scrambled inverted faces), performance is significantly worse than if only three are (as in the new Thatcherized inverted faces). This makes it all the more remarkable that the upright new Thatcherized faces (three features inverted) are numerically superior to the upright scrambled faces (all six features upright). It is this result that makes possible the main effect of stimulus type, indicating that overall performance is superior to that for new Thatcherized faces.

The last experiment (Experiment 3) in this article aimed to adduce additional evidence for the inversion effect in the new Thatcherized faces, confirm the elimination of the FIE found with the 50% feature-inverted and scrambled faces reported in Civile, McLaren, et al.'s (2014) studies, and investigate whether there is an orientation **<x>** face type interaction for new Thatcherized faces compared with 50% feature-inverted and scrambled faces. These latter faces are, in effect, scrambled new Thatcherized faces, and so this last experiment combines both the transformations used separately in Experiment 2 in this condition.

EXPERIMENT 3

The purpose of this experiment was to compare the FIE obtained for categories of 50% feature-inverted and scrambled faces used in Civile, McLaren, et al.'s (2014) experiments with that obtained for sets of new Thatcherized faces used in Experiments 1 and 2 in this article.

METHOD

Materials

Experiment 3 adopted once again the four categories of scrambled faces used in Experiment 2, but to generate the 50% feature-inverted and scrambled faces, we turned half the features interior to each face upside down. Specifically, for each of the four category prototypes, we inverted one of the eyes, one of the ears, and either the nose or the mouth. As was the case in the previous experiment, each scrambled face drawn from a given category had the location and orientation of its features specified by its category prototype. These new stimuli now had half of their features inverted and half in their usual upright orientation (i.e., they were 50% feature-inverted and scrambled faces).

Participants

A total of 32 psychology undergraduates at the University of Exeter took part in the experiment. We counterbalanced the study, as in Experiment 1 and 2, by splitting the participants into eight groups.

Procedure

The procedure was exactly the same as that used in Experiments 1 and 2.

RESULTS

The mean latencies for the four stimulus conditions were 50% feature-inverted and scrambled faces upright, 2.65 s; 50% feature-inverted and scrambled faces inverted, 2.70 s; new Thatcherized upright, 2.35 s; and new Thatcherized inverted, 2.42 s. The data from all 32 participants were used in the signal detection d' analysis (Figure 5). An ANOVA revealed a significant main effect of face type, $F(1, 31) = 4.090$, MSE $= 0.313$, $p = .049$, $d = 0.44$, 95% CI [0.30, 0.59]. There was a trend that approached significance for the main effect of orientation, $F(1, 31) = 3.228$, MSE $= 0.385$, $p = .07$, $d = 0.42$, 95% CI [0.27, 0.57]. Finally, there was a

significant interaction between face type and orientation, $\underline{F}(1, 31) = 4.619$, $\underline{MSE} = 0.035$, $\underline{p} =$ .047, $\underline{d} = 0.51$, 95% CI [0.23, 0.78]. Thus, simple effect analyses were conducted, showing an inversion effect for new Thatcherized faces, $\underline{t}(31) = 3.150$, $\underline{SE} = 0.123$, $\underline{p} = .003$, $\underline{d} = 0.74$, 95% CI [0.57, 0.92], but there was no effect of inversion for 50% feature-inverted and scrambled faces, $\underline{t}(31) = 0.102$, $\underline{SE} = 0.154$, $\underline{p} = .91$. Performance in recognizing new Thatcherized upright faces was significantly better than recognition for upright 50% feature-inverted and scrambled faces, $\underline{t}(31) = 3.098$, $\underline{SE} = 0.126$, $\underline{p} = .004$, $\underline{d} = 0.70$, 95% CI [0.50, 0.89], and also significantly better than recognition for inverted 50% feature-inverted and scrambled faces, $\underline{t}(31) = 2.242$, $\underline{SE} = 0.167$, $\underline{p} = .032$, $\underline{d} = 0.62$, 95% CI [0.41, 0.83]. Additional analyses showed that all conditions were recognized significantly above chance level, with smallest $\underline{t}$ for inverted 50% feature-inverted and scrambled faces, $\underline{t}(31) = 2.165$, $\underline{SE} = 0.119$, $\underline{p} = .038$, $\underline{d} = 0.54$, 95% CI [0.37, 0.70].

<Insert Figure 5 about here>

DISCUSSION

The results of Experiment 3 confirmed the results of our earlier experiments. The new Thatcherized faces this time produced a significant inversion effect, and the size of this effect was significantly greater than that for 50% feature-inverted and scrambled faces. The 50% feature-inverted and scrambled faces showed no sign of any inversion effect while keeping recognition performance above chance levels (i.e., agreeing with Civile, McLaren, et al.'s 2014 finding). In effect, this demonstrates that the significant inversion effect found in this experiment with new Thatcherized faces can be eliminated by scrambling the features of these faces. It is also noteworthy that three of the four stimulus conditions presented in Experiment 3 showed a very similar performance level. These three conditions were matched in terms of the number of facial features presented in upright and in inverted orientations (all 50%). Two of the conditions,

17

those involving 50% feature-inverted and scrambled faces, had all their configural information disrupted by our manipulations, whereas the new Thatcherized inverted faces still had first-order configural information similar to that in normal (inverted) faces. Thus, this seems to suggest that inversion prevents the use of this first-order information. The advantage participants had in recognizing upright new Thatcherized faces (which also had 50% of their features inverted) makes a strong case for the first-order configural information still present in this face driving the inversion effect with these stimuli. For this to be true we would have to assume that this type of information is effective only when the face is in its normal orientation.

GENERAL DISCUSSION

In the three experiments reported in this article we investigated whether configural information affects the FIE. Experiment 1 produced equivocal results, in that a type of Thatcherization manipulation did not entirely eliminate the effect of inversion while reducing it significantly. A trend toward a significant inversion effect for the new Thatcherized faces was found in Experiment 2, and this provided good evidence for the effect when taken in combination with Experiment 1. Experiment 3 confirmed this effect. We can now extend our Bayesian analysis to incorporate the results of Experiment 3, which on its own has a Bayes factor of 53. The combined Bayes factor for all three experiments is thus just over 297, which establishes this result beyond reasonable doubt. The results of these studies suggest that if we control for the influence of local feature orientation but leave first-order configurations unaltered, we can obtain an inversion effect with these stimuli. And finally, we were able to confirm that the inversion effect disappears only when both the configural information and local feature orientation information are manipulated so as to no longer be differentially available in the upright orientation. Given that the difference between the two sets of faces in Experiment 3

18

was whether they were scrambled or not, the significant interaction here allows us to conclusively answer the question "Does configural information play a part in the FIE?" by saying "yes." The main finding of these experiments is that they provide clear evidence that it is possible to obtain an inversion effect when configural information is the main factor in play. This raises the question of why this type of information would be one of the factors that helps face discrimination in the upright condition.

If one thinks first about first-order information, it is what is commonly shared across every facial exemplar, in some sense the configuration corresponding to the average of the spatial relationships between the features inside a face. A potential explanation for its influence on the FIE comes from studies supporting the holistic process account of face recognition. Mondloch and Maurer (2008) suggested that holistic processing has been tuned to upright faces. In their study on the composite face effect (CFE) they found that holistic processing decreased linearly over the entire range of orientations but remained significant when faces were oriented at 30° or 60° from upright. When faces reached a sideways orientation (90°), the CFE was present in the means but no longer statistically significant, and with further rotation it disappeared altogether. The results are especially convincing because the diminution of holistic processing was revealed by increased accuracy on "same" trials as the faces were rotated further from the upright. This is contrary to the usual decrease in accuracy of face processing with rotation (Collishaw & Hole, 2002; Mondloch, Le Grand, & Maurer, 2002; Valentine & Bruce, 1988) but just as would be predicted if the holistic processing that makes these trials hard for upright faces were diminishing. Their results extend the many previous reports that the CFE seen for upright faces is not present for inverted faces (Carey & Diamond, 1994; Hole, 1994; Young, Hellawell, & Hay, 1987).

We can start to explain our findings, then, by proposing that first-order configural information presented in a familiar (upright) orientation is used to engage holistic processing and that this confers some advantage in face recognition. This analysis finds some support in the work of Maurer, Le Grand, and Mondloch (2002), who described the term configural processing as constituting three different processes: sensitivity to first-order relational information, sensitivity to second-order relational information, and holistic processing (i.e., perceiving the detected features as a whole, making it harder to process individual features). If this position were taken, then it could be said that the new Thatcherized stimuli used in our experiments contain considerable first-order configural information and that this configural information might be expected to lead to the face being processed holistically in such a way that individual features play less of a role. Thus, the upright new Thatcherized faces are processed holistically and benefit from this processing, and the orientation of individual features within them matters less. When inverted these faces are not processed holistically, and the orientation of individual features becomes more important, which leads to performance equivalent to our 50% feature-inverted and scrambled faces in either orientation.

In line with this view, Rossion (2008) referred to holistic processing as a perceptual process and not a cue that can be manipulated on a single face. Thus, the term configural indicates physical information that can be measured and disrupted in a stimulus (by changing the spatial distance between features), whereas holistic face processing seems to refer to "the simultaneous integration of several features of a face into a single perceptual representation"[B] (Rossion, 2008). The consequence of this integration is the fact that the facial features are interdependent, so the observer cannot focus on one feature of the face stimulus without being influenced by the other features at the same time. Likewise, our results seem to suggest that, as

long as enough of the first-order relations are kept unaltered in manipulated faces, in a recognition task individuals will still benefit from the holistic representation of those stimuli.

Furthermore, Hole, George, and Dunsmore (1999) suggested that there is more than one type of relational processing. Thus, they interpreted their results (that recognition of the top half of a composite face, constructed from top and bottom halves of different faces, is difficult when the face is upright but not when it is inverted, even for image negatives) by suggesting that upright negative chimeric faces are sufficiently facelike to evoke holistic processing. In this view, holistic processing is elicited by anything that roughly conforms to the basic plan of a face, and it is holistic encoding that establishes that it is a face that is being perceived, as opposed to some other kind of object. By contrast, configural processing deals with the precise locations of the facial features relative to one another. According to Hole et al. (1999), it may be that inversion disrupts both holistic and configurational processing, whereas constructing a photographic negative of the facial image disrupts configural processing but leaves holistic processing intact. All this is consistent with the idea that it is first-order relations that elicit the facelike perception of the Thatcherized faces, leading to the holistic analysis that benefits discrimination in the upright orientation. And perhaps it is this holistic processing that explains the higher performance participants showed for new Thatcherized faces compared with scrambled ones in Experiment 2, producing the main effect of face type in that experiment. Thus, this suggests that first-order relations provide a vital substrate needed to activate a holistic process for perception of facelike stimuli. Of course, we can only infer this from our results, because we did not directly test for holistic processing in our experiments, but this explanation does seem to fit both our and other researchers' data well.

If instead we allowed a more robust version of Hole et al.'s (1999) holistic processing construct, one that would imply that our sets of scrambled faces could be identified as faces and so generate an inversion effect, we would then have to explain why this holistic processing ceased to apply to the 50% feature-inverted and scrambled faces. If the explanation was that a certain number or proportion of facial features had to be upright for holistic processing to apply and only the upright scrambled faces met this criterion, then we would be left in some difficulty in explaining why performance on the 50% feature-inverted and scrambled faces is superior to that on the inverted scrambled faces; surely they should be the same. Given that this is not the case, it makes more sense to argue that the strong holistic construct suggested by Hole et al. (1999) is elicited only when a face shows at least some evidence of the configural information we expect in a face, in this case first-order relational information in an upright orientation.

What can we say about the role of second-order relational information in our experiments? Our new Thatcherized stimuli necessarily disrupt this second-order information by changing the detailed relations between stimulus features, even though the basic layout of the face has not been greatly altered. Because this manipulation is fundamentally more disruptive to second-order information than first-order, we have suggested the first-order information as the configural component that drives the effect by engaging holistic processing of the face when it is upright. Nevertheless, some second-order information is intact (that between the features that have not been inverted), and this could also be playing a role. As configural information, it could simply contribute to engaging holistic processing, but this type of configural information is capable of doing more than that.

Although first-order configural information is very useful for identifying a face as a face, it is not so useful for discriminating between faces. Thus, the advantage contingent on holistic

face is not likely to stem from first-order configural information, and almost by definition it cannot come from local feature information, but it could be due to second-order configural information because this will differentiate between different faces. The full explanation of our results would then appeal to both types of configural information. First-order effects allow the engagement of a holistic processing mode that emphasizes second-order configural information relative to local features. Expertise in the use of this information then allows superior recognition performance for upright faces (see Civile, Zhao, et al., 2014, for a recent article arguing for the role of expertise). This advantage can be removed by inverting the face, because this prevents holistic processing, or by scrambling the face, which changes the configural information available so that our expertise in dealing with it is no longer applicable. Performance is then determined by local feature processing, which is better dealt with in the upright (i.e., familiar) orientation. Thus, we are arguing that the existence of some second-order relational information driving the inversion effect with the new Thatcherized stimuli cannot be ruled out, and clearly this would be consistent with Diamond and Carey's (1986) original position. More research on the role of first- and second-order relational information and its link with holistic processing in bringing about the FIE is needed to establish whether this account is correct.

In conclusion, taken together the results reported in this article confirm that it is possible to obtain an inversion effect when all the configural information in a face is changed, but they have also shown that it is possible to obtain an inversion effect when the local feature information is controlled for while the first-order relations are kept as close as possible to normal. Thus, the new Thatcherized faces used in Experiments 1, 2, and 3 in combination establish that an inversion effect can be obtained with those stimuli. The implication of this

finding is that there seem to be two sources of information affecting facial recognition contingent

on inversion, local feature orientation, and the configural relationships between those features.

References

Bartlett, J., & Searcy, J. (1993). Inversion and configuration of faces. Cognitive Psychology, 25, 281-316.

Carey, S. (1992). Becoming a face expert. Philosophical Transactions Royal Society of London, 335, 395-403.

Carey, S., & Diamond, R. (1994). Are faces perceived as configurations more by adults than by children? Visual Cognition, 1, 253-274.

Carey, S., Diamond, R., & Woods, B. (1980). The development of face recognition: A maturational component? Developmental Psychology, 16, 257-269.

Civile, C., McLaren, R., & McLaren, I. P. L. (2014). The face inversion effect: Parts and wholes. Quarterly Journal of Experimental Psychology, 67, 728-746.

Civile, C., Zhao, D., Ku, Y., Elchlepp, H., Lavric, A., & McLaren, I. P. L. (2014). Perceptual learning and inversion effects: Recognition of prototype-defined familiar checkerboards. Journal of Experimental Psychology: Animal Learning and Cognition, 40, 144-161.

Collishaw, S., & Hole, G. (2002). Is there a linear or a nonlinear relationship between rotation and configural processing of faces? Perception, 31, 287-296.

Diamond, R., & Carey, S. (1986). Why faces are and are not special: An effect of expertise. Journal of Experimental Psychology: General, 115, 107-117.

Dienes, Z. (2011). Bayesian versus orthodox statistics: Which side are you on? Perspectives on Psychological Science, 6, 274-290.

Hole, G. (1994). Configurational factors in the perception of unfamiliar faces. Perception, 23, 65-74.

Hole, G., George, P., & Dunsmore, V. (1999). Evidence for holistic processing of faces viewed as photographic negatives. Perception, 28, 341-359.

Leder, H., & Bruce, V. (1998). Local and relational aspects of face distinctiveness. Quarterly Journal of Experimental Psychology, 51, 449-473.

Leder, H., & Bruce, V. (2000). When inverted faces are recognized: The role of configural information in face recognition. Quarterly Journal of Experimental Psychology, 53, 513-536.

Maurer, D., Le Grand, R., & Mondloch, J. (2002). The many faces of configural processing. Trends in Cognitive Sciences, 6, 255-260.

McKone, E., & Yovel, G. (2009). Why does picture-plane inversion sometimes dissociate perception of features and spacing in faces, and sometimes not? Toward a new theory of holistic processing. Psychonomic Bulletin and Review, 16, 778-797.

Mondloch, C., Le Grand, R., & Maurer, D. (2002). Configural face processing develops more slowly than featural face processing. Perception, 31, 553-566.

Mondloch, J., & Maurer, D. (2008). The effect of face orientation on holistic processing. Perception, 37, 1175-1186.

Rakover, S., & Teucher, B. (1997). Facial inversion effects: Parts and whole relationship. Perception & Psychophysics, 59, 752-761.

Rhodes, G., & Tremewan, T. (1994). Understanding face recognition: Caricature effects, inversion and the homogeneity problem. Visual Cognition, 1, 275-311.

Rossion, B. (2008). Picture-plane inversion leads to qualitative changes in face perception. Acta Psychologica, 128, 274-289.

Searcy, J., & Bartlett, J. (1996). Inversion and processing of component and spatial-relational

    information in faces. Journal of Experimental Psychology: Human Perception and

    Performance, 22, 904-915.

Tanaka, J., & Farah, M. (1991). Second-order relational properties and the inversion effect:

    Testing a theory of face perception. Perception & Psychophysics, 50, 367-372.

Thompson, P. (1980). Margaret Thatcher: A new illusion. Perception, 9, 483-484.

Valentine, T. (1988). Upside down faces: A review of the effect of inversion upon face

    recognition. British Journal of Psychology, 79, 471-491.

Valentine, T., & Bruce, V. (1988). Mental rotation of faces. Memory & Cognition, 16, 556-566.

Yin, R. (1969). Looking at upside-down faces. Journal of Experimental Psychology, 81, 141-

    145.

Young, A., Hellawell, D., & Hay, D. (1987). Configurational information in face perception.

    Perception, 16, 747-759.

<Captions start here>

Figure 1. Examples of stimuli used in Experiment 1 showing the 4 category prototypes for the new Thatcherized faces in both orientations and an example of the normal faces. The dimensions of the stimuli were 5.63 cm **x** 7.84 cm, and the stimuli were presented at a resolution of 1,280 **x** 960 pixels

Figure 2. Mean $\underline{d}'$ for each of the 4 facial conditions in the old/new recognition phase, Experiment 1

Figure 3. Examples of the category prototypes for the scrambled faces used in Experiment 2 and previously used in the Civile et al. (2014) study

Figure 4. Mean $\underline{d}'$ for each of the 4 facial conditions in the old/new recognition phase, Experiment 2. Dashed lines specify comparisons and $\underline{p}$ values

Figure 5. Mean $\underline{d}'$ for each of the 4 facial conditions in the old/new recognition phase, Experiment 3