


This article was downloaded by: [Carnegie Mellon University], [Michael J. Tarr]
On: 21 February 2012, At: 12:33
Publisher: Psychology Press
Informa Ltd Registered in England and Wales Registered Number: 1072954
Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH,
UK



Visual Cognition

Publication details, including instructions for authors and subscription information:

<http://www.tandfonline.com/loi/pvis20>

Recognizing disguised faces

Giulia Righi ^a, Jessie J. Peissig ^b & Michael J. Tarr ^c

^a Children's Hospital Boston, Harvard Medical School, Boston, MA, USA

^b Department of Psychology, California State University Fullerton, Fullerton, CA, USA

^c Department of Psychology, Carnegie Mellon University, Pittsburgh, PA, USA

Available online: 13 Feb 2012

To cite this article: Giulia Righi, Jessie J. Peissig & Michael J. Tarr (2012): Recognizing disguised faces, *Visual Cognition*, 20:2, 143-169

To link to this article: <http://dx.doi.org/10.1080/13506285.2012.654624>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.tandfonline.com/page/terms-and-conditions>

This article may be used for research, teaching, and private study purposes. Any substantial or systematic reproduction, redistribution, reselling, loan, sub-licensing, systematic supply, or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae, and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand, or costs or damages

whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

Recognizing disguised faces

Giulia Righi¹, Jessie J. Peissig², and Michael J. Tarr³

¹Children's Hospital Boston, Harvard Medical School, Boston, MA, USA

²Department of Psychology, California State University Fullerton, Fullerton, CA, USA

³Department of Psychology, Carnegie Mellon University, Pittsburgh, PA, USA

Across three experiments, we evaluated the effects of “disguises” on observers’ face identification performance using naturalistic images in which individuals posed with a variety of wigs and eyeglasses. Experiment 1 tested recognition memory performance with and without disguises and showed that any changes in these facial attributes hindered performance, replicating previous findings. More interestingly, Experiment 1 revealed that a change in hairstyle or the removal of eyeglasses had more impact on performance than did the addition of eyeglasses. In Experiment 2, disguised and undisguised faces were presented upright or inverted to examine whether the performance decrements seen in Experiment 1 were attributable to disruption of processing strategies for faces. Despite showing an overall effect of disguise on performance, there was no interaction between inversion and disguise. This suggests that disguises are not directly affecting processing strategies, but instead disguises are likely encoded as part of the overall representation. Similarly, in Experiment 3, the composite face paradigm was used to examine whether the impact of disguises is attributable to disruption of holistic processing for faces. Again, we found an overall effect of disguise on performance, but here we also observed that the presence of disguises interacted with participants’ tendency to form holistic face representations.

Keywords: Disguise; Face inversion; Face recognition; Object categorization; Perceptual expertise.

Please address all correspondence to Jessie J. Peissig, Box 6846, Department of Psychology, California State University Fullerton, Fullerton, CA 92834, USA. E-mail: jpeissig@fullerton.edu

This research was funded by NSF Award No. 0339122 (Enhancing Human Performance), the Perceptual Expertise Network (No. 15573-S6), a collaborative award from James S. McDonnell Foundation, and by the Temporal Dynamics of Learning Center at UCSD (NSF Science of Learning Center SBE-0542013). We would like to thank Michele Hebert and Yi Cheng for their assistance in data collection.

© 2012 Psychology Press, an imprint of the Taylor & Francis Group, an Informa business
<http://www.psypress.com/viscog> <http://dx.doi.org/10.1080/13506285.2012.654624>

“God has given you one face, and you make yourself another”

(William Shakespeare)

The ability to recognize thousands of people by their faces is one of the hallmarks of the human experience. This uncanny ability has interested psychologists for decades and the result is a very extensive literature on both the human proficiency at face recognition and its limitations. For example, it has been shown that recognition performance is relatively impervious to physical transformations including blurring (Harmon & Julesz, 1973a, 1973b; Yip & Sinha, 2002), changes in lighting conditions (Braje, 2003; Braje, Kersten, Tarr, & Troje, 1998), and changes in viewing angle (Hill, Schyns, & Akamatsu, 1997; O’Toole, Edelman, & Bülthoff, 1998). Our resilience to these transformations is something that we all have experienced. At the same time, there are some facial transformations to which we are particularly sensitive. For example, it may take us a moment to recognize someone we see every day if they suddenly shave off their beard or dramatically shorten their hair. Such recognition difficulty may be even more pronounced if one is trying to recognize a relatively unfamiliar person who has changed their appearance. Intentional transformations can dramatically disguise one’s identity and can include a wide variety of altered physical attributes, including wearing a wig, changing hairstyle or hair colour, wearing eyeglasses, removing or growing a beard, etc.

Despite the frequency of appearance changes in our real life experience (even omitting intentional “disguises”), few psychological studies have systematically studied how such appearance changes affect face identification (but see Diamond & Carey, 1977; Metzger, 1999; Patterson & Baddeley, 1977; Terry, 1994). Nonetheless, the consistent result across this handful of previous studies is that a disguise—almost by definition—hinders face recognition in both children (Diamond & Carey, 1977; Metzger, 1999) and adults (Patterson & Baddeley, 1977; Terry, 1993, 1994). Of note, Patterson and Baddeley (1977) found that changing either hairstyle or beard decreased facial recognition memory performance equally compared to undisguised faces, and that changes in both attributes produced chance-level performance. Subsequently, Terry (1994) found that facial recognition memory was most impaired by the removal of eyeglasses or by the addition or removal of a beard. Such results provide concrete evidence for a lack of invariance in face identification with changes in some facial attributes. However, it is an open question as to why this is so.

One possible explanation is it that all featural alterations impair face recognition equally, as suggested by the results of Patterson and Baddeley (1977). Alternatively, it is possible that different types of disguises, which either change a relatively stable facial feature (i.e., hairstyle) or an attribute that is often added or subtracted (i.e., eyeglasses), affect face recognition

differently. Moreover, the size of the effect of the disguise could change depending on whether it is present when a face is first studied or if it is introduced in subsequent presentations, as suggested by Terry (1994). These are the issues that will be addressed in Experiment 1, by testing recognition memory performance across a variety of disguise manipulations appearing on natural faces.

The results of prior experiments using disguises also raise questions regarding how disguises interfere with the processing mechanisms used for face recognition. One of the most salient findings of the face processing literature is that faces are not identified by focusing on specific features in isolation, but rather that observers rely on both specific facial features and the relationship among these features, performing what has been termed “holistic” processing (Farah, Wilson, Drain, & Tanaka, 1998; Hole, 1994; Tanaka & Farah, 1993; Young, Hellawell, & Hay, 1987). More recently, several authors have suggested that there are different types of holistic processing that come into play in face recognition, and that can be specifically tested using different manipulations (Gauthier & Tarr, 2002; Maurer, Le Grand, & Mondloch, 2002; Rotshtein, Geng, Driver, & Dolan, 2007). Differences in terminology aside, both Gauthier and Tarr (2002) and Maurer and colleagues (2002) suggest that when viewing faces, observers process: (1) The specific arrangement of features across all faces (i.e., two eyes, above nose and mouth); (2) the specific spatial relations between features in a given face (i.e., distance between the eyes, length of the forehead, etc.); and (3) a face gestalt in which all facial features are integrated into a single representation.

Taken together, these processing strategies have been implicated as the major reason as to why face identification is relatively robust across many physical transformations such as changes in viewpoint, luminance, and blurring (e.g., Braje et al., 1998; Hole, George, Eaves, & Rasek, 2002; O’Toole et al., 1998). Thus, given both anecdotal and empirical evidence that the same cannot be said for disguises, we wondered whether face identification is susceptible to disguises because they affect an observer’s ability to apply one or more of these processing strategies. It is this question that will be addressed in Experiments 2 and 3.

In Experiment 2 we examine whether disguises interact with specific processing strategies of faces. For example, a change in hairstyle might produce a change in the perceived distance between the eye region and the hairline, or the addition or removal of eyeglasses might alter the perceived distance between the eyes. This will be examined using the well-documented face inversion effect (Yin, 1969), which assesses identity recognition across picture-plane inversion. Consider two alternative outcomes. Recent studies have shown that inversion disrupts both configural and part-based processing. Thus, if disguises lead to a disruption in configural processing and/or

part processing (similar to documented effects caused by inversion; McKone & Yovel, 2009), we might expect a selective impairment in the processing of upright disguised faces. Under this view, there would be smaller difference between upright and inverted faces with disguises, and a *smaller* inversion effect for disguised faces compared to undisguised faces. Alternatively, observers may encode varying disguised features as invariant facial features regardless of the recognition strategies deployed. That is, disguises may generally disrupt recognition in that some facial features have changed in appearance. Under this view, we would expect an overall decrease in performance for disguised faces across both the upright and inverted conditions.

In Experiment 3 we examine whether disguises interact with the default observer bias to form a gestalt or holistic face representation. For example, the obvious presence of disguises might bias observers towards using more stable, local facial features that are not affected by the disguises, thus disrupting our default tendency to form holistic face representations. We address this question using the composite face paradigm (Hole, 1994; Young et al., 1987), which provides a measure of the bias to fuse different face parts into a single representation. In this paradigm composite faces are created by combining the top part of one face and the bottom part of another face to create a new face; participants are then asked to specifically recognize only one of the two parts. The holistic processing bias is derived from the fact that interference arises from the presence of the irrelevant face half. If disguises interact with such holistic processing, we would expect a reduction in the tendency of our participants to fuse face halves into a unitary representation.

EXPERIMENT 1

Method

Participants. Participants were 20 Brown University students. All participants provided informed consent and received \$7 or course credit for their participation.

Design and materials. Experiment 1 consisted of two phases: A learning phase and a testing phase. In the learning phase, participants were sequentially shown a series of faces and asked to rate them along several physical dimensions to promote familiarization with the stimuli without explicitly suggesting memorization. After the learning phase, participants performed an old-new recognition memory task on a series of faces that included all of the faces in the learning set plus an equal number of foil, never-before-seen faces.

Face stimuli were taken from the TarrLab face database (www.face-place.org). This database contains photographs of many different individuals in various types of disguises, such that, for each individual, there are multiple photographs in which hairstyle and/or eyeglasses have been changed/added. All faces were cropped to include the outer contour of the face, hairstyle and a small portion of the neck of each individual (Figure 1). The faces were also normalized in size to a standard interocular distance.

In the learning phase, the face stimulus set was comprised of 48 faces (24 females and 24 males). As shown in Table 1, for each sex, 12 of the faces were shown with no disguise (e.g., natural hairstyle and no eyeglasses). The remaining 12 faces were shown disguised with either wig only (four faces); eyeglasses only (four faces); or a wig and eyeglasses (four faces). Each individual face was presented in only one version during learning.

In the testing phase, the face stimulus set was comprised of 96 faces, of which 48 were new to the participant and 48 were familiar to the participant from the learning phase. For each sex, the new faces included 12 disguised and 12 undisguised faces. Of the 12 disguised faces, four were presented with a wig only, four with eyeglasses only, and four with both a wig and eyeglasses (identical to the learning phase; Column 1 of Table 1). The familiar faces were presented either in the same version that participants had seen during learning (12 total for each sex: Six disguised and six undisguised), or in a different version compared to the learning phase (also 12 for each sex). Of




Figure 1. Experiments 1–3, sample face stimuli drawn from the TarrLab face database (www.face-place.org; stimulus images courtesy of Michael J. Tarr, Center for the Neural Basis of Cognition, Carnegie Mellon University, <http://www.tarrlab.org>), with the original face (left) and the disguised version of the same individual (right). Note that in Experiments 1 and 2 these stimuli were presented in intact form; in Experiment 3 top and bottom halves of faces were combined to create new “composite” face images. To view this figure in colour, please see the online issue of the Journal.

TABLE 1
Experiment 1 procedure in the learning and testing phases for one sex

<i>Learning phase</i>	<i>Testing phase</i>
Wig only Disguised = 4 (with a wig)	Wig only Trained disguised Tested disguised = 2 (with wig, no change) Tested undisguised = 2 (wig subtracted, 1 change)
Undisguised = 4 (natural hair)	Trained undisguised Tested disguised = 2 (with wig added, 1 change) Tested undisguised = 2 (natural hair, no change)
Glasses only Disguised = 4 (with eyeglasses)	Glasses only Trained disguised Tested disguised = 2 (with eyeglasses, no change) Tested undisguised = 2 (glasses subtracted, 1 change)
Undisguised = 4 (no eyeglasses)	Trained undisguised Tested disguised = 2 (with eyeglasses added, 1 change) Tested undisguised = 2 (no eyeglasses, no change)
Wig + Eyeglasses Disguised = 4 (with eyeglasses + wig)	Wig + Eyeglasses Trained disguised Tested disguised = 2 (with eyeglasses + wig, no change) Tested undisguised = 2 (glasses + wig subtracted, 2 changes)
Undisguised = 4 (no eyeglasses + natural hair)	Trained undisguised Tested disguised = 2 (glasses + wig added, 2 changes) Tested undisguised = 2 (no eyeglasses + wig, no change)

the faces that were changed from study to test, eight of each sex were shown in a version that differed only in one attribute from the learning phase, that is either the hairstyle or eyeglasses were different, and four for each sex were shown in a version that differed in two attributes from the learning phase, that is both eyeglasses and hairstyle were changed (Column 2 of Table 1).

Procedure. In this and all subsequent experiments, stimuli were presented using an Apple iMac G5 20-inch computer (Apple Inc., Cupertino, CA) and the Psychophysics Toolbox for Matlab (Mathworks, Natick, MA; Brainard, 1997; Pelli, 1997).

The learning phase consisted of three runs, each of which included all 48 training faces. In the first run participants judged whether they thought each face was attractive, in the second run they judged whether each face was male or female, and in the third run they judged whether each face was younger or older than 21 years of age. Each participant took a self-timed break between each judgement task, and then began the next task at their own pace. Each face was displayed until the participant made a response, with an intertrial interval (ITI) of 500 ms.

The testing phase consisted of a single run which included all of the individual faces shown during the training phase, as well as an equal number of new, never-before-seen individual faces for a total of 96 trials. On each trial, the participant decided whether they remembered or did not remember seeing that particular individual's face during the learning phase ("old/new" task). Participants were explicitly instructed to ignore any changes in facial attributes and any disguise manipulation. Each face was displayed until the participant made a response using one of two keys corresponding to old and new.

Results

Several analyses were carried out in order to address the experimental questions. First, individual participants' recognition memory performance was assessed by computing d' (sensitivity for discriminating old from new items) separately for 0, 1, and 2 attribute changes from learning to test. One change corresponded to the hair being changed relative to the face shown during the learning phase (from a wig to natural, or natural to a wig), or to eyeglasses being either added or subtracted relative to each training face. Two changes corresponded to the hair being changed in conjunction with eyeglasses being either added or subtracted. Figure 2 shows mean d' values across participants for each level of change. Planned pairwise comparisons were performed between no-change and 1-change and 1-change and 2-change trials in order to establish whether performance worsens as more attributes are changed between learning and test. These comparisons revealed a significant difference between no-change and 1-change, $t(19) = 9.16$, $p < .0001$, two-tailed; Cohen's $d = 1.21$, and a significant difference between 1-change and 2-change, $t(19) = 3.6$, $p < .002$, two-tailed; Cohen's $d = 0.71$.

Second, we analysed a subset of the 1-change trials in order to examine whether the types of changes across stimulus presentations (addition vs. subtraction vs. change of a disguise) differentially affected recognition

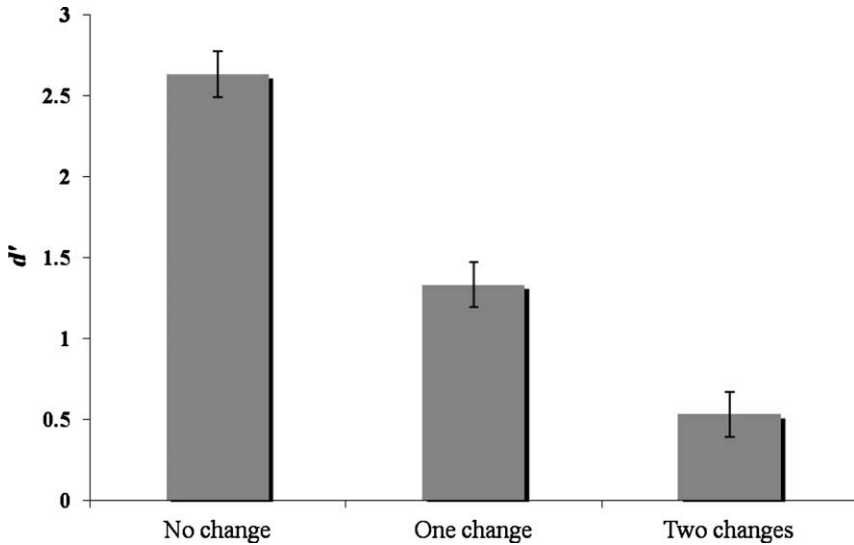


Figure 2. Experiment 1, d' for three conditions. No change: Faces seen in identical manner at study and test (mean = 2.63, $SD = 1.18$); one change: Faces that changed hairstyle or added/subtracted eyeglasses between study and test (mean = 1.33, $SD = 0.96$); two changes: Faces that changed hairstyle and added/subtracted eyeglasses between study and test (mean = 0.53, $SD = 1.28$).

performance. For each subject we computed average percent correct for trials in which the hairstyle was changed from study to test, trials in which eyeglasses were added between study and test, and trials in which eyeglasses were subtracted between study and test. Planned comparisons were performed between the addition and subtraction of eyeglasses, between a change in hairstyle and the addition of eyeglasses, and between a change in hairstyle and the subtraction of eyeglasses (Figure 3). The first pairwise t -test revealed a significant difference between adding eyeglasses versus subtracting eyeglasses, $t(19) = 3.2$, $p < .005$, two-tailed; Cohen's $d = 0.99$, showing that participants were significantly more accurate when a disguise attribute was introduced at test, as compared to when the same disguise attribute was removed at test (Figure 3). The second pairwise t -test revealed a significant difference between changing hairstyle and adding eyeglasses, $t(19) = 4.22$, $p < .0005$; Cohen's $d = 0.95$, with better performance when eyeglasses were added as compared to when hairstyle was changed. The third pairwise t -test revealed no significant difference between changing hairstyle and subtracting eyeglasses, $t(19) = 1.14$, $p = 0.27$, two-tailed. Taken together these comparisons demonstrate that the addition of a disguise between study and test is the least detrimental manipulation, whereas changes in features that are present during the study phase are somewhat more damaging to face recognition.




Figure 3. Experiment 1, average percentage correct for a subset of trials in which eyeglasses were added between study and test (mean = 0.675, $SD = 0.20$), a subset of trials in which eyeglasses were subtracted between study and test (mean = 0.41, $SD = 0.32$), and a subset of trials in which hairstyle was changed between study and test (mean = 0.49, $SD = 0.19$).

Discussion

Experiment 1 replicated findings from several previous studies (Diamond & Carey, 1997; Patterson & Baddeley, 1977; Terry, 1994), in that it revealed a significant decrease in recognition memory performance when disguise attributes are changed between study and test. Thus, we can confidently state that face identification is not invariant over disguises. Beyond this validation, Experiment 1 adds to our understanding of the impact of disguises in the following ways.

First, our present results indicate that the magnitude of the disguise effect is dependent on the number of changes that occur between when a face is first learned and when it is later presented. We posit that when observers learn a new face they encode a representation that is specific, at least in part, to local attributes such as eyeglasses and/or hairstyle. Local changes in these specific attributes individually influence later identification. In particular, as the number of changes increases there is increasing dissimilarity between the encoded version of an individual and the disguised version of that individual.

Second, our results indicate that changes in relatively stable facial features such as hairstyle or the removal of eyeglasses exert a greater cost on

performance than does the addition of disguise elements. This greater cost of subtraction as compared to addition is consistent with the results previously reported by Terry (1994), who suggested that a face with eyeglasses is less distinctive than a face without eyeglasses in the context of other faces with eyeglasses (i.e., people with eyeglasses all look more similar) and as a consequence a face that was studied without eyeglasses will be more easily recognizable because of its relative distinctiveness. However, this explanation does not seem to fit our present experiment, as eyeglasses were neither extremely rare nor very common. More specifically, in Experiment 1 study faces with eyeglasses did not represent the majority of faces, and the proportion of faces with and without eyeglasses was not different between study and test; thus a change in the relative distinctiveness of the study faces seems unlikely to account for our findings.

An alternative account of the results of Experiment 1 is that specific facial attributes (e.g., a specific hairstyle and/or eyeglasses) become integral to a more global face representation. Although more weight is typically assigned to the internal features of a face, Hancock, Bruce, and Burton (2000) have shown that unfamiliar faces are encoded, at least in part, using external facial features such as the specific hairstyle. As such, when observers are presented with unfamiliar faces, changing the hairstyle or removing eyeglasses creates a greater mismatch with respect to the overall face representation compared to the addition of eyeglasses, and this mismatch is not easily compensated for without significant alteration of that face representation. The finding that removal of a feature is more detrimental than addition may be part of a more general visual encoding strategy, and not specific to face processing per se. Whatever features are present at encoding automatically become part of the representation of that object (including surface properties such as texture; Vuong, Peissig, Harrison, & Tarr, 2005); thus, removal of the feature negatively impacts recognition. Later adding a feature to an object will have an impact, but far less so than removal.

It is also possible that the presence of a disguise produces an asymmetric cost of modifying the perceived configural properties of a face between encoding and recognition. That is, both a different hairstyle and eyeglasses could affect the processing of spatial relations, or second-order relations, between the features of a face (Maurer et al., 2002). For example, a different hairstyle could change the appearance of a person's hairline and thus modify the perceived distance between the eyes and the hair. As such, a change in hairstyle or encoding a face with eyeglasses might produce a representation of the face that is configurally different from the studied version, and, as a consequence, less identifiable. Alternatively, a face encoded without eyeglasses could give rise to a more stable representation because its configural properties are encoded without the presence of any occluder. Hence, even

with the subsequent addition of eyeglasses, faces encoded without eyeglasses might be more easily identified given a more robust preexisting configural representation.

Finally, a combination of these explanations is also possible. A change in hairstyle may hinder performance because hairstyle is encoded as part of the face representation, whereas a change in the presence or absence of eyeglasses may disrupt the perception of the configural properties of a face. Unfortunately, Experiment 1 does not allow us to disambiguate these different hypotheses.

In summary, our present results both confirm prior work demonstrating that disguises make face identification more difficult, and add to the literature regarding the patterns of increased difficulty we see with different levels of disguise. Interestingly, we also observe asymmetries in recognition performance depending on the both the content of the change and whether the change was an addition or subtraction of a facial feature. As discussed earlier, there are several explanations for this asymmetry. Do disguises affect recognition simply because of the change in the face between study and test, or, alternatively, do they affect recognition because they disrupt/modify second-order spatial relations and/or parts processing? In that the recognition memory paradigm used in Experiment 1 does not allow us to distinguish between these alternatives, we designed Experiments 2 and 3 to more precisely examine what components of the face recognition process are disrupted by different aspects of disguises.

EXPERIMENT 2

Our goal with Experiments 2 and 3 was to better understand which of several factors—alone or in combination—underlie the significant decrement in facial recognition performance we observed with disguises. For example, disguise elements might interact with the perception of spatial/configural properties of a given face—for example, a wig changing the perceived distance between the eyes and the hairline. Disguise elements may also interact with the specific part-based processing strategies, thereby differentially altering recognition performance. Finally, observers may simply (and with a potential cost) encode changeable attributes within face representations.

Picture-plane inversion has been previously assumed to disrupt the processing of second-order relations in faces (e.g., variation in the spacing between features; for a review, see Maurer et al., 2002). However, more recent reviews have questioned the interpretation of inversion as disrupting only second-order relations (McKone & Yovel, 2009; Riesenhuber & Wolff, 2009; Rossion, 2008, 2009). For example, McKone and Yovel (2009) looked at feature and spacing information in the context of upright and inverted

faces to determine whether inversion effects were driven primarily by spacing information or if feature information also contributes to inversion effects. In a review of 17 different studies, they found the expected large inversion effects for spacing. McKone and Yovel also found inversion effects for feature information, and the size of these feature inversion effects was often large. Thus, inversion effects may be the result of spacing or feature changes. Consequently, our inversion manipulation won't allow us to discern feature changes from spatial changes, but it will allow us to compare the inversion effect for disguised and undisguised faces. If we find a significant interaction, then this result indicates that disguises are fundamentally changing the processing of the features and/or spacing in faces (specifics can then be determined later). If we do not find a significant interaction, then this result indicates that disguises are not fundamentally changing the way faces are processed. Instead, disguises are just decreasing performance overall, perhaps by the encoding of changeable features, such as glasses and hair/wigs, as part of the original face representation.

On a given trial in Experiment 2, participants saw a pair of sequentially presented faces that both appeared in either the upright or the inverted orientation.¹ Participants were instructed to decide whether the two faces were of the same individual or not, irrespective of the presence or absence of disguise attributes. That is, they were explicitly told to make a same-different judgement based on face identity and not on face attributes, such as hairstyle or eyeglasses.

Method

Participants. Participants were 20 Brown University students. All participants provided informed consent and received \$7 or course credit for their participation. Data from one participant was lost due to equipment malfunction; thus, the analyses reported here include only 19 participants.

Design and materials. The face stimuli were again drawn from the TarrLab face database. Stimuli were chosen to be disguised with either changes in hairstyle, changes in eyeglasses, or changes in hairstyle and eyeglasses together. Each face stimulus was normalized to have the same interocular distance.

Experiment 2 consisted of 320 trials. One half of the trials (160) were presented in the upright orientation, and the other half (160) were presented in the inverted orientation. Within each group of these trials, 80 trials showed faces with the same identity (same trials), and the other 80 showed

¹ Studies of face inversion generally use a similar design in which both comparison faces appear at the same orientation.

faces of different identities (different trials). Eighty different identities (40 males and 40 females) were presented four times each. For both the upright and the inverted orientations the same-identity trials consisted of: 80 no-disguise trials (the same image across the two presentations), 40 1-change trials (26 faces had a hairstyle change; 14 faces had eyeglasses added), and 40 2-change trials (40 faces with a hairstyle change and eyeglasses added). For both the upright and the inverted orientations the different-identity trials consisted of: 84 trials in which both identity and hairstyle was different, and 76 trials in which faces differed in identity, hairstyle, and in the presence of eyeglasses. The slight mismatch in the number of faces across disguise conditions was due to the variable availability of specific disguise configurations within the face database used. Priority was given to repeating each individual identity an equal number of times (for a schematic breakdown of all different trial types, see Table 2). Note that if a trial included eyeglasses (for both same and different trials), the image with eyeglasses was always shown second. Trial types were presented in random order.

Procedure. We were unable to include trials with glasses subtracted as well as glasses added due to a lack of available stimuli. Participants performed a two-alternative forced choice (2AFC) sequential-matching task. On each trial the first face was centrally presented for 250 ms, followed by a mask presented for 500 ms, and then the second face centrally presented for 250 ms, again followed by a mask for 500 ms. The mask consisted of 10×10 pixel squares taken from the set of faces used in the experiment placed randomly within a canvas that matched the dimensions of the face

TABLE 2
Distribution of trials across experimental conditions in Experiment 2

<i>Same trials</i>		
No disguise = 80	Upright = 40 Inverted = 40	
One change = 40	Hairstyle change = 26 Glasses added = 14	Upright = 13 Inverted = 13 Upright = 7 Inverted = 7
Two changes = 40	Upright = 20 Inverted = 20	
Different trials		
Different individuals = 84	Upright = 44 Inverted = 44	
Different individuals with eyeglasses = 76	Upright = 38 Inverted = 38	

images (480 × 400 pixels). Participants judged whether or not the two face images belonged to the same individual, regardless of possible changes in facial appearance. Judgements of face identity were used because we wanted participants to focus on the entire face and not any particular local facial feature such as hairstyle or eyeglasses. Responses and reaction times were recorded from the onset of the second face with no time limit on responding. Participants responded using one of two keys corresponding to same and different. The ITI following the participant's response was 500 ms.

Results

Performance in the sequential-matching task was measured by computing percentage correct for “same trials” separately for the upright and the inverted orientations across four levels of disguise (no change, hairstyle change, eyeglasses added, hairstyle changed and eyeglasses added; for group means, see Figure 4). A repeated measures two-way ANOVA was computed using percentage correct as the dependent measure across two factors: Disguise (four levels) and orientation (two levels). This ANOVA revealed a main effect of disguise, $F(3, 18) = 20.7$, $p < .0001$; Cohen's $d = 1.45$, suggesting overall poorer matching performance when a disguise was present, and a main effect of orientation, $F(1, 18) = 12.6$, $p < .003$; Cohen's $d = 0.6$, suggesting overall poorer matching performance with inverted faces. No significant interaction was found, $F(3, 54) = 0.63$, $p = .6$, suggesting that the magnitude of the inversion effect was not significantly different despite the presence or absence of a disguise. In order to examine the differential impact of the three levels of disguises (i.e., hairstyle change, eyeglasses added, hairstyle change and eyeglasses added), several post hoc Tukey HSD tests were carried out. These tests revealed that at the $p < .05$ threshold there was no significant difference between changes in hairstyle and the addition of eyeglasses, $p > .05$; Cohen's $d = 0.17$, but that both of these manipulations produced significantly better performance as compared to a hairstyle change combined with the addition of eyeglasses: Glasses added vs. hairstyle/eyeglasses change, $p < .05$; Cohen's $d = 0.55$; hairstyle change vs. hairstyle/eyeglasses, $p < .05$; Cohen's $d = 0.66$.

Discussion

Consistent with the results of Experiment 1, Experiment 2 revealed a significant main effect of disguise, indicating that changes in facial attributes negatively impact recognition performance. This effect of disguise holds even in instances where the disguise manipulation is readily apparent to participants and their memory load is minimal. These findings offer further evidence that observers *do not* encode invariant face representations, at least

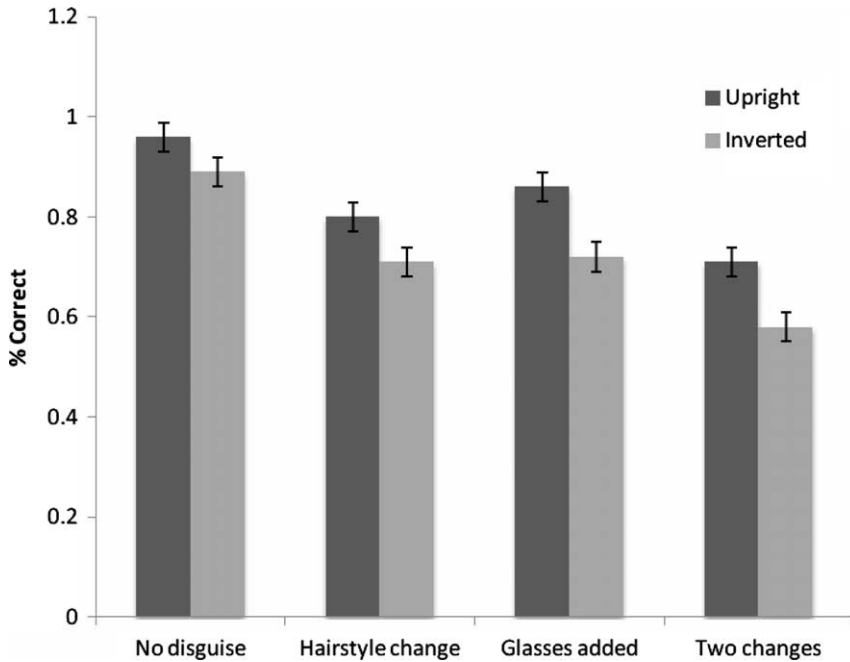


Figure 4. Experiment 2, average percentage correct for the sequential-matching task for: upright faces no disguise (mean = 0.96, $SD = 0.05$) and inverted faces no disguise (mean = 0.89, $SD = 0.09$); upright faces with a hairstyle change (mean = 0.79, $SD = 0.17$) and inverted faces with a hairstyle change (mean = 0.71, $SD = 0.2$); upright faces with eyeglasses added (mean = 0.85, $SD = 0.16$) and inverted faces with eyeglasses added (mean = 0.72, $SD = 0.29$); upright faces with hairstyle change and eyeglasses added (mean = 0.71, $SD = 0.15$) and inverted faces with hairstyle change and eyeglasses added (mean = 0.58, $SD = 0.25$).

with respect to changes in facial attributes that nominally would not be held to affect identity processing. Moreover, the short-term nature of the task in Experiment 2 suggests that the detrimental effect of disguises found in Experiment 1 was not driven by the high memory load imposed by learning 48 individual faces. In contrast to the results of Experiment 1, performance was not differentially affected by a change in hairstyle and the addition of eyeglasses, as both of these disguises produced a similar decrement in performance. One possible explanation for this result is that, in this experiment, participants did not have the chance to form a relatively stable face representation in that they saw each face very briefly. Under such conditions any change in physical appearance may be treated similarly, producing indistinguishable performance decrements.

With respect to the primary goal of Experiment 2—gaining a better understanding as to why disguises affect identity recognition—this question

was addressed by comparing the performance on upright and inverted faces across disguise conditions. In Experiment 2, no interaction was found between disguises and orientation, suggesting that the cost of inversion did not differ across presentations of undisguised and disguised faces, nor was it modulated by the type of disguise applied to a face. This result suggests that the presence of a disguise does not affect the face recognition strategy of faces. Finally, as can be seen in Figure 4, inversion affected all three disguise conditions equally, suggesting that different sorts of superficial facial changes have about the same effect on recognition performance.

To summarize, the results of Experiment 2 demonstrate that the presence of disguises affects identity recognition even when there is no significant long-term memory load. More importantly, the presence of disguises is not interacting with a part-based processing strategy, nor is it interacting with the ability to perceive spatial distances between features. Thus, the primary effect of disguises on identity recognition may be due to the fact that disguise attributes, such as hairstyle or eyeglasses, are encoded as part of the overall representation of an individual's face despite their relative instability. However, Experiment 2 does not allow us to entirely exclude the possibility that observers may be modifying processing strategy between disguised and undisguised faces. Namely, it is possible that the presence of disguises interferes with a second strategy that is automatically applied to faces, namely the bias to fuse all facial features into a unitary, holistic representation. This holistic processing has been proposed as a type of configural processing distinct from the processing of second-order relations (Maurer et al., 2002). The increased perceptual processing load created by the introduction of a disguise could push observers to focus on a more localized region of each face, thereby decreasing one's typical bias to process faces as gestalts. This possibility is examined in Experiment 3.

EXPERIMENT 3

Experiment 3 was designed to examine whether disguises interact with holistic processing of faces using the "composite face" paradigm (Hole, 1994; Young et al., 1987). In a "complete" composite face paradigm participants are asked to match either the top part (i.e., from the nose up) or the bottom part (i.e., from the nose down) of a face across sequential presentations (Cheung, Richler, Palmeri, & Gauthier, 2008). The relevant parts can be either the same or different, and so can the irrelevant parts. This gives rise to congruent trials (i.e., trials in which both relevant and irrelevant parts lead to the same answer) and incongruent trials (i.e., trials in which relevant and irrelevant parts lead to different answers). Holistic processing bias can be measured by comparing performance between congruent and

incongruent trials; differential performance across these two conditions shows that participants have difficulty ignoring the irrelevant face part, and thus have a propensity to fuse both the relevant and irrelevant face parts into a unitary representation (Richler, Tanaka, Brown, & Gauthier, 2008). In Experiment 3 we adapted the composite face paradigm to include disguised faces. Disguises took the form of changes in hairstyle or eyeglasses and were present only on the top part of a face, although participants were asked to perform judgements on both the top and bottom face parts.

Method

Participants. Participants were 19 Brown University students. All participants provided informed consent and received \$7 or course credit for their participation.

Design and materials. The face stimuli were again drawn from the TarrLab face database. Stimulus preparation for this experiment consisted of dividing each face into two halves, such that the top and bottom parts could be combined across individuals. Nine male and nine female faces were chosen, making sure to avoid any hairstyle that would extend to the lower part of the face, thereby making it obvious that the two parts did not belong to the same individual. The distribution of disguises was as follows: Six female faces were used in their original version and with eyeglasses added,² and three female faces were used in their original version and with a hairstyle change; six male faces were used in their original version and with a hairstyle change, and three male faces were used in their original version and with eyeglasses added. All disguised versions were selected so as to have one change in appearance (i.e., addition of eyeglasses or a new hairstyle) compared to their undisguised version. In order to minimize differences in width and make our composites more realistic, the stimuli were also normalized in size using the outer contour of each face.

The experiment consisted of 144 trials divided into eight experimental conditions (see Figure 5), such that there were 18 trials in each condition (nine trials original and nine trials disguised). In half of the trials participants were presented with faces that had no disguise (original; Figure 5, top). These trials were distributed among four conditions: Face tops and face bottoms could both be the same across sequential presentation (same and congruent), face tops and bottoms could both be different across sequential presentations (different and congruent), face tops could be the same while the bottoms were different (same and incongruent for a top

²We were unable to include trials with glasses subtracted as well as glasses added due to a lack of available stimuli.

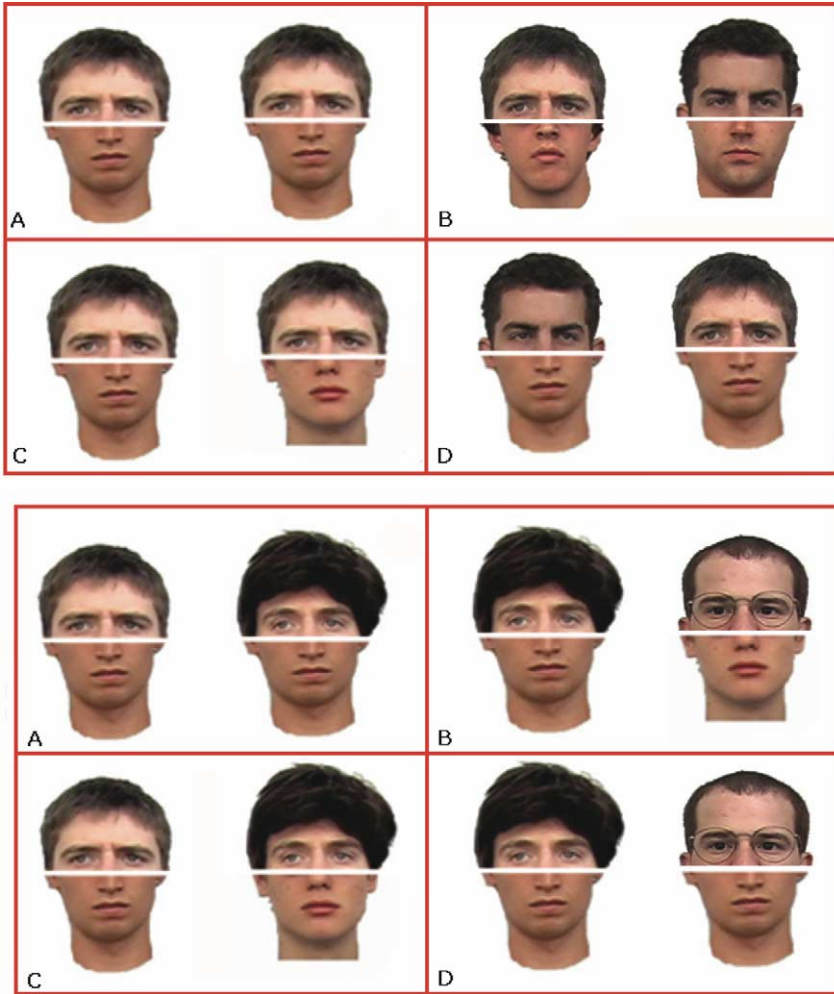


Figure 5. (Top) Experiment 3, examples of original (undisguised) experimental trials for top and bottom judgements. (A) Same congruent (top and bottom); (B) different congruent (top and bottom); (C) same incongruent (top) or different incongruent (bottom); (D) different incongruent (top) or same incongruent (bottom). (Bottom) Experiment 3, examples of disguised experimental trials for top and bottom judgements. (A) Same congruent (top and bottom) with hair change; (B) different congruent (top and bottom) 1 with eyeglasses; (C) same incongruent (top) or different incongruent (bottom) with hair change; (D) different incongruent (top) or same incongruent (bottom) with eyeglasses. To view this figure in colour, please see the online issue of the Journal.

judgement, different & incongruent for a bottom judgement), and lastly face tops could be different while the bottoms were the same (different and incongruent for a top judgement, same and incongruent for a bottom

judgement). In the other half of the trials participants were presented with faces that displayed a disguise created by either changing hairstyle or by adding eyeglasses (Figure 5, bottom). These two types of disguises appeared an equal number of times. These trials were also distributed among four conditions: Face tops and bottoms could be the same but have a hairstyle change or eyeglasses added across sequential presentation (same and congruent), face tops and bottoms could both be different and eyeglasses would be present in one of the two sequential presentations (different and congruent), face tops could be the same but with a disguise present, while bottoms were different (same and incongruent for a top judgement, different and incongruent for a bottom judgement), and lastly face tops could be different also in the presence of eyeglasses while bottoms are the same (different and incongruent for a top judgement, same and incongruent for a bottom judgement). Different individuals *always* differed in hairstyle. All eight trial types appeared an equal number of times throughout the experiment in a random order.

Procedure. Participants performed a sequential-matching task on either the tops or the bottoms of two “composite” faces. Each trial began with a 250 ms cue informing the participant whether he/she had to match the tops or the bottoms. The first composite face was then presented for 500 ms, followed by a mask identical to the one used in Experiment 2 that was also presented for 500 ms. Finally, a second composite face was presented for 500 ms. The intertrial interval (ITI) was 500 ms. Participants were instructed to be as fast and as accurate as possible in making their responses.

Results

Performance for each participant was assessed by computing d' measures of accuracy separately for congruent and incongruent trials across top judgements of original (undisguised) faces, top judgements of disguised faces, bottom judgements of original faces, and bottom judgements of disguised faces. Figure 6 shows mean d' values for all eight conditions. A three-way repeated measure ANOVA was computed using disguise (present, absent), trial type (congruent, incongruent), and part judged (top, bottom) as factors. This ANOVA revealed a significant main effect of disguise, $F(1, 18) = 22.6$, $p < .0002$; Cohen's $d = 0.52$, showing overall worse performance when disguises were present. There was also a significant main effect of trial type, $F(1, 18) = 15.6$, $p < .001$; Cohen's $d = 0.55$, indicating an overall cost of incongruency across all trial types. The d' did not significantly differ between top and bottom part judgements, $F(1, 18) = 0.52$, $p = .48$. The ANOVA also revealed two interactions. First, there was a significant Disguise \times Part judged interaction, $F(3, 54) = 48.5$, $p < .0001$, Cohen's $d = 1.6$, suggesting a

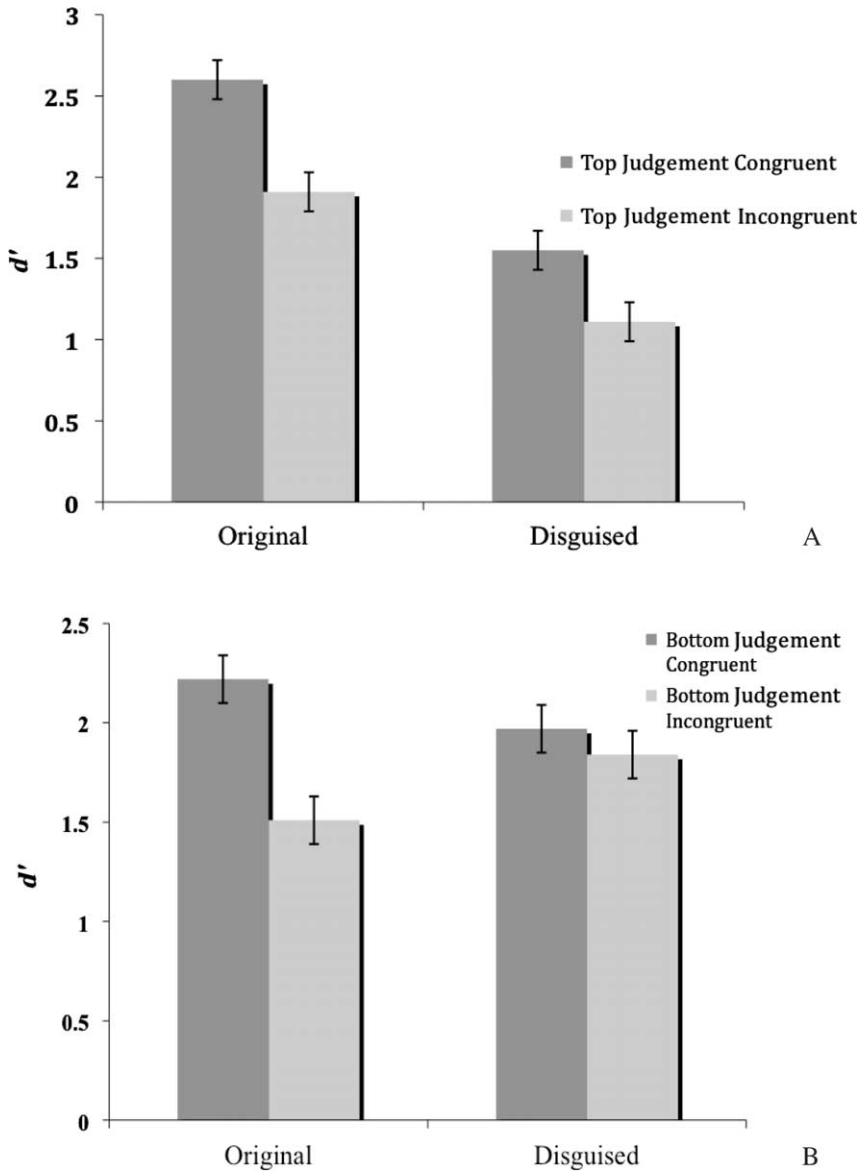


Figure 6. Experiment 3, (A) mean d' measure of accuracy for judgements on top face part: Congruent undisguised (mean = 2.59, $SD = 0.63$), incongruent undisguised (mean = 1.91, $SD = 1.28$), congruent disguised (mean = 1.55, $SD = 0.88$), incongruent disguised (mean = 1.11, $SD = 0.95$); (B) mean d' measure of accuracy for judgements in bottom face part: Congruent undisguised (mean = 2.22, $SD = 0.82$), incongruent undisguised (mean = 1.51, $SD = 1$), congruent disguised (mean = 1.97, $SD = 0.75$), incongruent disguise (mean = 1.84, $SD = 0.72$).

differential effect of disguises on top and bottom judgements. Tukey HSD tests were carried out in order to understand the nature of this interaction. These comparisons revealed that while judgements for top face parts were significantly affected by the presence of disguises, $p < .05$; Cohen's $d = 0.98$, there was no significant effect of disguise on bottom judgements, $p > .05$; Cohen's $d = 0.11$. Second, there was a marginally significant interaction between disguise and trial type, $F(3, 54) = 2.9$, $p = .1$, Cohen's $d = 0.41$. Tukey HSD tests revealed that this interaction was driven by a significant effect of congruency for original trials, $p < .05$; Cohen's $d = 0.72$, but no significant effect of congruency for disguised trials, $p > .05$; Cohen's $d = 0.35$. No significant three-way interaction was found, $F(1, 18) = 0.92$, $p = .35$.

Discussion

Experiment 3 investigated the holistic processing bias in the context of disguises with the goal of better understanding how changes in identity-irrelevant facial attributes affect processing strategies. First, the data showed a significant main effect of disguise, which provided further confirmation of the fact that the presence of a disguise hinders recognition performance.

Second, the data showed a main effect of congruency, showing better performance on congruent trials compared to incongruent trials. The logic behind the composite face paradigm being used to assess holistic processing is that there are two alternate conditions: On one hand the top and bottom of a face can lead to the correct answer (congruent); in contrast, the top and bottom halves can lead to two competing responses (incongruent). Thus if observers cannot help fusing the top and bottom halves (e.g., they automatically form a "holistic" representation of each face), in the incongruent trials they are more likely to provide an incorrect response because of the competition between the two possible responses. In this framework the effect of congruency on matching performance can be taken as a measure of the bias for creating a holistic representation of the two face parts fused into a single identity (Gauthier & Bukach, 2007; Richler et al., 2008). Thus our data suggest that participants are processing the composite faces holistically.

The results of Experiment 3 also revealed two interactions: A significant interaction between the face part judged and the presence of disguises, and a marginally significant interaction between disguise and congruency. The first interaction was driven by a significant effect of disguises on judgements of face tops and the absence of an effect of disguises on judgements of face bottoms. Because disguises only ever appeared on the top halves of faces, one could expect that disguises would be more detrimental to recognition for top judgements as compared to bottom judgements. However, holistic processing should render the disguises as potent regardless of where the

disguises appear and what part the participants are attending to. Our present results suggest that in the presence of disguises, observers may modify their processing strategy to be better able to attend to the relevant half and ignore the irrelevant half of each face. This conjecture is also supported by the marginally significant interaction between the size of the congruency effect and the presence of disguises, as this interaction is driven by the lack of significant effect of the congruency manipulation when disguises were present, in contrast to a significant effect of this manipulation when disguises were absent.

In sum, Experiment 3 suggests that the presence of disguises interacts with a holistic processing bias when participants are asked to process face bottoms, but not when participants are asked to process face tops. This may be the case because the disguises we used changed the appearance of the top halves, which may drive participants to look at the bottom halves to help them create a more “complete” face representation even in cases where the bottom halves should be ignored. In other words, the presence of a disguise may increase the perceptual weight given to unmodified face parts, which in the case of a top half judgement leads to an *increase* in holistic processing. Alternatively, when participants are asked to make judgements on the bottom halves, they might not be weighing information from the top halves as much as they would in the absence of a disguise, in that the disguise manipulations make it more difficult to extract meaningful information from the disguised part.

GENERAL DISCUSSION

Almost all of us are experts at face recognition (Duchaine & Nakayama, 2005; Tanaka, 2001), and our perceptual expertise with faces appears to be very robust to a wide variety of changes that we encounter in our day-to-day lives, such as changes in facial expression, lighting conditions or viewing angle. There is, however, one class of real-world facial transformations with which—at least anecdotally—most observers will have experienced some difficulty: Changes in the physical appearance of a face through alterations of hairstyle, adding or removing eyeglasses and beards, covering the hair, etc. Despite the prevalence of these types of transformations in real life, relatively few studies have systematically assessed their impact on face recognition (but see Diamond & Carey, 1977; Metzger, 1999; Patterson & Baddeley, 1977; Terry, 1993, 1994). The goal of our present study was to systematically examine what types of disguises are most damaging to face recognition, and also to explore how these transformations affect performance.

Consistent with earlier studies, Experiment 1 found that changes in physical appearance impair the recognition of unfamiliar faces. Moreover, by

using several types of disguises the present results showed that recognition memory decreases in relation to the number of changes between study and test—the more changes the poorer the performance. Experiment 1 also produced a relatively novel finding, showing that recognition was significantly more impaired when hairstyle was changed or eyeglasses were subtracted between study and test, compared to the addition of eyeglasses. This result is consistent with the work of Terry (1994), who focused on the impact of eyeglasses and beards on recognition. He proposed that the differential effect of eyeglasses between study and test could be attributed to differences in distinctiveness. He suggested that faces with eyeglasses are less distinct from one another, compared to faces with no disguise. As a consequence, when a face that was originally learned with eyeglasses is seen without eyeglasses, it might appear more different compared to other faces presented within that set, and thus might not be recognized easily (Terry, 1994). However, this explanation does not account for our current findings. First of all, in the present study faces with eyeglasses did not represent the majority of faces, making it unlikely that they would have influenced the perception of distinctiveness in the face set used. Moreover, performance was equally affected by a change in hairstyle and by the subtraction of eyeglasses, even though presumably the presence of very different hairstyles across the face set used in the present experiment should have maximized the diversity of the faces. The differential impact of these different disguises instead suggest that a specific hairstyle and eyeglasses are encoded as part of a relatively stable face representation for an individual during the learning phase; in contrast, when eyeglasses are presented only at test, observers are better able to ignore them and match the face with the representation without eyeglasses. As a consequence, learned faces that appear with a new hairstyle might look more different from the original representation, compared to a learned face that at study appears with eyeglasses.

Although Experiment 1 provides further evidence of the detrimental effect of disguises on identity recognition, it still fails to provide evidence regarding how processing strategies are disrupted when facial attributes are manipulated. There are at least two hypotheses that can be derived from the results of Experiment 1. First of all, the present results suggest that observers may rely on local facial features both during the encoding and the recognition of faces, such that when one or more of these features is disguised recognition becomes more difficult. The reliance on local features could be due to the use of unfamiliar faces, as previous studies have found that when presented with unfamiliar faces, participants tend to use an image-matching type strategy that weighs local features more heavily than configural cues (Bruce et al., 1999; Hancock et al., 2000; Megreya & Burton, 2006). Alternatively, it is possible that different types of disguises affect the second-order relations

within the face, or the ability to use specific strategies such as parts-based processing. For example, second-order relations may be disrupted by a change in hairstyle that changes the perceived vertical distance between the eyes relative to the hairline, or eyeglasses that might either interfere with the perception both of horizontal and vertical relations between the eye region and the edges of a face or act as occluders.

In order to test these hypotheses, Experiment 2 relied on the well-documented face inversion effect (Yin, 1969) in which picture-plane inversion of a face is used to disrupt processing strategies, both part-based and configural (McKone & Yovel, 2009). Based on work carried out by Goffaux and Rossion (2007), we hypothesized that if disguises affected the processing strategies of a face, we should observe a smaller cost of inversion for disguised faces compared to undisguised faces. However, if disguises acted primarily in the context of creating a less stable face representation, as suggested by the results of Experiment 1, we should only observe an overall cost of disguise on performance, but no interactions.

The results of Experiment 2 showed that the presence of a disguise hinders recognition performance even when faces are presented sequentially, but they failed to show a differential cost of inversion across changes in facial attributes. Thus, Experiment 2 suggests that changing local facial attributes, whether it is a hairstyle change, the introduction of eyeglasses, or both, does not affect either perception of second-order relations or parts-based processing.

Finally, regarding the possible mechanisms that underlie the decrement in performance observed in the presence of disguises, it is possible that the obvious presence of changes in physical attributes might influence another configural processing strategy applied to face perception, namely the bias to fuse different face parts and features into a unitary representation (Gauthier & Tarr, 2002; Maurer et al., 2002). For example, the increased difficulty introduced by disguises could bias observers to attend more to local facial features and, in turn, this could reduce the ability to create a face gestalt. Experiment 3 was designed to measure the impact of disguises on holistic processing using the composite face paradigm (Hole, 1994; Young et al., 1987). In the classic composite face paradigm, participants are presented with faces created by merging the top and bottom of two different people and they are asked to recognize or match either the top or the bottom half of each face across stimulus presentations. Studies have shown that observers process composite faces in a manner that fuses the different individual parts into a new identity, thereby making it difficult to accurately judge an individual part when, for example, face parts are presented one on top of the other (e.g., Hole, 1994; Young et al., 1987). In the case of disguises, composites can be used to assess whether the manipulation of specific features (i.e., hairstyle, adding eyeglasses) affects the perception of only the

face part in which disguises are present (the top), or whether it also interacts with the global face configuration by influencing the perception of both parts (top and bottom).

Experiment 3 replicated the findings of Experiment 1 and 2 with regard to overall cost of disguise on identity recognition performance. However, the results also suggest that the presence of a disguise interacts with the specific judgements participants were asked to make. The results revealed a significant interaction of disguise (disguised vs. undisguised) by part judged (top vs. bottom) and a marginally significant interaction of disguise (disguised vs. undisguised) by trial type (congruent vs. incongruent). Such results point to the fact that although judgements to the top halves were directly impaired in the presence of disguises, judgements to the bottom halves were not. It is reasonable to believe that, given that disguises render face recognition more challenging, in the bottom half condition participants might be trying harder to focus their attention on the relevant face part (the bottom), and thus were less affected by the type of top part presented. In contrast, for top half judgements, congruency has a large effect. One possible explanation is that because disguises are changing the appearance of the top halves, participants are incorrectly (at least for incongruent trials) looking at the bottom halves to help them create a more “complete” or more “stable” face representation. Thus, in the case of top half judgements participants’ recognition processes are biased towards relying more on holistic processing rather than the parts. This suggests that disguises affect both local feature processing and holistic face perception. Taken together, these results suggest that people can learn to adjust their weighting of featural or holistic cues depending on the specific demands imposed by the recognition context. Thus, if you have a coworker who frequently changes hairstyle, you may be able to change the weight your facial recognition system places on that feature for that individual.

In summary, the present experiments explored how changes in facial attributes—disguises—affect face identification. The main goal of these experiments was to understand why these real-world changes in facial appearance are surprisingly difficult to deal with. Across all three experiments, data showed that observers were worse at recognizing faces that had been disguised in some manner, and, moreover, that the impact of this manipulation is dependent on the type of disguises used. However, the results of Experiment 3 suggest that the introduction of a disguise might interact with the bias to create a holistic face percept. It is reasonable to believe that disguises on unfamiliar faces are difficult to deal with because they become part of the encoded face representation, and because they interact with the holistic processing bias. Thus, we may be able to train observers to be less affected by disguises and to ignore hairstyle and eyeglasses while forming a face representation, at the same time maintaining

holistic processing strategies. Such training might enable security and law enforcement personnel to better ignore features such as eyeglasses and hairstyle, making them more accurate at recognizing the individuals they are tasked with identifying.

REFERENCES

- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433–436.
- Braje, W. L. (2003). Illumination encoding in face recognition: Effect of position shift. *Journal of Vision, 3*(2), 161–170.
- Braje, W. L., Kersten, D., Tarr, M. J., & Troje, N. F. (1998). Illumination effects in face recognition. *Psychobiology, 26*(4), 371–380.
- Bruce, V., Henderson, Z., Greenwood, K., Hancock, P. J. B., Burton, A. M., & Miller, P. I. (1999). Verification of face identities from images captured on video. *Journal of Experimental Psychology: Applied, 5*(4), 339–360.
- Cheung, O. S., Richler, J. J., Palmeri, T. J., & Gauthier, I. (2008). Revisiting the role of spatial frequencies in the holistic processing of faces. *Journal of Experimental Psychology: Human Perception and Performance, 34*, 1327–1336.
- Diamond, R., & Carey, S. (1977). Developmental changes in the representation of faces. *Journal of Experimental Child Psychology, 23*, 1–22.
- Duchaine, B., & Nakayama, K. (2005). Dissociations of face and object recognition in developmental prosopagnosia. *Journal of Cognitive Neuroscience, 17*(2), 249–261.
- Farah, M. J., Wilson, K. D., Drain, M., & Tanaka, J. W. (1998). What is “special” about face perception? *Psychological Review, 105*(3), 482–498.
- Gauthier, I., & Bukach, C. (2007). Should we reject the expertise hypothesis? *Cognition, 103*(2), 322–330.
- Gauthier, I., & Tarr, M. J. (2002). Unraveling mechanisms for expert object recognition: Bridging brain activity and behavior. *Journal of Experimental Psychology: Human Perception and Performance, 28*(2), 431–446.
- Goffaux, V., & Rossion, B. (2007). Face inversion disproportionately impairs the perception of vertical but not horizontal relations between features. *Journal of Experimental Psychology, Human Perception and Performance, 33*(4), 995–1002.
- Hancock, P. J. B., Bruce, V., & Burton, A. M. (2000). Recognition of unfamiliar faces. *Trends in Cognitive Science, 4*, 330–337.
- Harmon, L. D., & Julesz, B. (1973a). Masking in visual recognition: Effects of two-dimensional filtered noise. *Science, 180*, 1194–1197.
- Harmon, L. D., & Julesz, B. (1973b). The recognition of faces. *Scientific American, 229*(5), 71–82.
- Hill, H., Schyns, P. G., & Akamatsu, S. (1997). Information and viewpoint dependence in face recognition. *Cognition, 62*(2), 201–222.
- Hole, G. J. (1994). Configurational factors in the perception of unfamiliar faces. *Perception, 23*(1), 65–74.
- Hole, G. J., George, P. A., Eaves, K., & Rasek, A. (2002). Effects of geometric distortions on face-recognition performance. *Perception, 31*(10), 1221–1240.
- Maurer, D., Le Grand, R., & Mondloch, C. (2002). The many faces of configural processing. *Trends in Cognitive Science, 6*(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*(5), 778–797.

- Megreya, A. M., & Burton, A. M. (2006). Unfamiliar faces aren't faces: Evidence from a matching task. *Memory and Cognition*, *34*(4), 865–876.
- Metzger, M. M. (1999). Differential effects of transformation on facial recognition in young children: A pilot study. *Perceptual and Motor Skills*, *89*(3), 799–807.
- O'Toole, A. J., Edelman, S., & Bühlhoff, H. H. (1998). Stimulus-specific effects in face recognition over changes in viewpoint. *Vision Research*, *38*(15–16), 2351–2363.
- Patterson, K. E., & Baddeley, A. D. (1977). When face recognition fails. *Journal of Experimental Psychology: Human Learning and Memory*, *3*, 406–417.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Richer, J. J., Tanaka, J. W., Brown, D. D., & Gauthier, I. (2008). Why does selective attention to parts fail in face processing? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(6), 1356–1368.
- Riesenhuber, M., & Wolff, B. (2009). Task effects, performance levels, features, configurations, and holistic face processing: A reply to Rossion. *Acta Psychologica*, *132*, 286–292.
- Rossion, B. (2008). Picture-plane inversion leads to qualitative changes of face perception. *Acta Psychologica*, *128*, 274–289.
- Rossion, B. (2009). Distinguishing the cause and consequence of face inversion: The perceptual field hypothesis. *Acta Psychologica*, *132*, 300–312.
- Rotshtein, P., Geng, J. J., Driver, J., & Dolan, R. J. (2007). Role of features and second-order spatial relations in face discrimination, face recognition, and individual face skills: Behavioral and functional magnetic resonance imaging data. *Journal of Cognitive Neuroscience*, *19*(9), 1435–1452.
- Tanaka, J. W. (2001). The entry point of face recognition: Evidence for face expertise. *Journal of Experimental Psychology: General*, *130*(3), 534–543.
- Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology*, *46A*(2), 225–245.
- Terry, R. L. (1993). How wearing eyeglasses affects facial recognition. *Current Psychology*, *12*(2), 151–162.
- Terry, R. L. (1994). Effects of facial transformations on accuracy of recognition. *Journal of Social Psychology*, *134*(4), 483–492.
- Vuong, Q. C., Peissig, J. J., Harrison, M. C., & Tarr, M. J. (2005). The role of surface pigmentation for recognition revealed by contrast reversal in faces and objects. *Vision Research*, *45*, 1213–1223.
- Yin, R. K. (1969). Looking at upside down faces. *Journal of Experimental Psychology*, *81*, 141–145.
- Yip, A. W., & Sinha, P. (2002). Contribution of color to face recognition. *Perception*, *31*, 995–1003.
- Young, A. W., Hellawell, D., & Hay, D. C. (1987). Configurational information in face perception. *Perception*, *16*, 747–759.