Journal of Experimental Social Psychology 74 (2018) 1–7

Contents lists available at ScienceDirect

Journal of Experimental Social Psychology

journal homepage: www.elsevier.com/locate/jesp

Case Report

Configural processing and social judgments: Face inversion particularly disrupts inferences of human-relevant traits☆

John Paul Wilson[a,*], Steven G. Young[b], Nicholas O. Rule[c], Kurt Hugenberg[d]

[a] Montclair State University, Department of Psychology, 1 Normal Ave, Montclair, NJ 07043, United States
[b] Baruch College & The Graduate Center, CUNY, Box B8-215, 55 Lexington Ave, New York, NY 10010, United States
[c] University of Toronto, Department of Psychology, 100 St. George Street, Toronto, ON M5S 3G3, Canada
[d] Miami University, 90 N Patterson Ave, Oxford, OH 45056, United States

ARTICLE INFO

Keywords:
Trait inferences
Person perception
Configural processing

ABSTRACT

Perceivers tend to strongly agree about the basic trait information that they encode from faces. Although some research has found significant consistency for social inferences from faces viewed at multiple angles, disrupting configural processing can substantially alter the traits attributed to faces. Here, we reconciled these findings by examining how disruptions to configural processing (via face inversion) selectively impairs trait inferences from faces. Across four studies (including a pre-registered replication), we found that inverting faces disrupted inferences about particularly human-relevant traits (trustworthiness and humanness) more than it did for a trait relevant to both human and non-human animals (dominance). These findings contribute to emerging research linking configural processing to the humanization of social targets, helping to provide a clearer understanding of how visual cognition may moderate perceptions of humanness.

Perceivers evaluate faces on basic social traits like trustworthiness and dominance quickly, efficiently, and with great consistency (e.g., Bar, Neta, & Linz, 2006; Willis & Todorov, 2006). Such evaluations can also have a strong impact: predicting the outcome of political elections (Hehman, Carpinella, Johnson, Leitner, & Freeman, 2014), economic decisions (van't Wout & Sanfey, 2008), and even life-and-death criminal sentences (Wilson & Rule, 2015, 2016). The cognitive and perceptual processes that people use to extract traits from faces remain a matter of some debate, however.

1. Featural versus configural processing

Central to the literature on face perception, *featural processing* and *configural processing* represent two fundamentally different ways of encoding faces. Faces have both features (e.g., eyes, nose, mouth) and a configuration (the eyes-over-nose-over-mouth arrangement typical of faces). When perceivers process faces featurally, they encode specific facial features to identify someone (such as a prominent nose or notable birthmark) without integrating features into a gestalt (Madera & Hebl, 2012; Tanaka & Farah, 1993). When perceivers process faces configurally, however, they extract and integrate information about multiple features in parallel and integrate them into a single representation that includes how the features relate to each other (see Maurer, Le Grand, & Mondloch, 2002, for a review).

These two processing modes can be distinguished using manipulations that interfere with the perception of configural versus featural information. For example, inverting a face (i.e., turning it upside-down) disrupts its eyes-over-nose-over-mouth configuration but not the perception of its constituent features. Because face inversion distinctly undermines configural processing, it has been used to explore what aspects of face perception specifically depend on configural information. Indeed, researchers commonly employ face inversion to demonstrate how configural processing influences a variety of outcomes, ranging from face memory (Yin, 1969) to recognizing facial expressions (Young & Hugenberg, 2010).

Recent evidence also suggests that face inversion can actually encourage dehumanization (e.g., Fincher & Tetlock, 2016). Hugenberg et al. (2016) found that face inversion slowed the processing of human-related concepts, disrupted categorizations of human faces as human, and reduced the ascription of human-like traits (e.g., the capacity for

☆ This research was supported by NSF grant BCS-1423765 and a grant from the Social Sciences and Humanities Research Council of Canada.
* Corresponding author.
E-mail addresses: wilsonjoh@montclair.edu (J.P. Wilson), Steven.Young@baruch.cuny.edu (S.G. Young), rule@psych.utoronto.ca (N.O. Rule), hugenbk@miamioh.edu (K. Hugenberg).

http://dx.doi.org/10.1016/j.jesp.2017.07.007
Received 29 April 2016; Received in revised form 2 May 2017; Accepted 20 July 2017
0022-1031/

emotional and cognitive sophistication). Put simply, the typical eyes-over-nose-over-mouth configuration of faces appears to serve as a bottom-up signal of humanness. This configural-humanness link runs in the opposite direction as well: Fincher and Tetlock (2016) found that the faces of dehumanized people were processed less configurally. Moreover, people tend to rely on featural processing when perceiving nonhuman faces (Dahl, Rasch, & Chen, 2014; Mondloch, Maurer, & Ahola, 2006). Thus, we propose that inferences of uniquely human traits might rely more on configural processing than do inferences of traits that humans share with other animals. To inform this hypothesis, we turn next to a brief description of how people infer traits from others' faces.

2. Inferring traits from faces

Perceivers extract trait information from faces both easily and with surprising consensus. For instance, Zebrowitz, Montepare, and Lee (1993) observed strong inter-rater consensus for a number of personality traits (warmth, dominance, strength, honesty, shrewdness) across faces of multiple ethnicities, and even young children make reliable inferences from faces following very brief exposures (e.g., 39 ms; Bar et al., 2006; Cogsdill, Todorov, Spelke, & Banaji, 2014). Data-driven models of social inferences have found that two orthogonal dimensions (facial *trustworthiness* and facial *dominance*) capture much of the variance in these consensual evaluations, and overlap strongly with facial expressions of emotion (Oosterhof & Todorov, 2008; Sutherland et al., 2013).

Consistent with our hypothesis, past studies have indirectly suggested that facial trustworthiness and dominance may rely on configural and featural information differently. In one example, participants judged facial halves as more trustworthy when paired with a trustworthy versus untrustworthy complementary half that they were instructed to ignore (Todorov, Loehr, & Oosterhof, 2010). Their seemingly involuntary integration of these “irrelevant” facial features with focal face characteristics into a unified percept suggests the importance of configural processing for extracting facial trustworthiness. Similarly, Hehman, Flake, and Freeman (2015) found that inferences related to intentions (i.e., trustworthiness) varied more across multiple presentations of the same face than did inferences related to ability (i.e., competence). Ability-related inferences typically relate to static and structural facial features whereas perceptions of trustworthiness typically depend on dynamic facial characteristics, such as facial affect (Carré, McCormick, & Mondloch, 2009; Hehman, Leitner, Deegan, & Gaertner, 2015; Oosterhof & Todorov, 2008). Thus, dynamic traits like trustworthiness seem to rely on constellations of features but structural traits like dominance may be gleaned from single cues (e.g., brow prominence, jaw size; Burton & Rule, 2013).

Given their foundational nature and potential processing distinction, we therefore focused on the “big two” traits of trustworthiness and dominance to examine the link between dehumanization and face perception. Specifically, we hypothesized that disrupting configural processing would affect inferences of especially human traits (e.g., trustworthiness) more than inferences of traits shared with animals (e.g., dominance). Indeed, although it may be difficult to imagine a cat or frog as trustworthy, animals' displays of dominant behavior tend to be quite clear. In fact, even *facial* dominance can be accurately observed in other mammals. Kramer, King, and Ward (2011) found that humans could accurately categorize chimpanzees' dominance by looking at their faces, but could not categorize them as sociable or sympathetic (i.e., traits often associated with humanness). We believe this may be due to different trait signals in faces: If features signal dominance but configurations signal sophisticated human-like traits (e.g., trustworthiness, empathy), then perceivers should be able to reliably extract dominance (but not sociability) from animal faces, as Kramer et al. (2011) found. No one has yet tested this, however.

3. Current research

Here, we therefore examined the role of configural processing in inferences of dominance and trustworthiness from people's faces. We predicted that inverting faces would disrupt the perception of traits considered uniquely human (e.g., trustworthiness) but not the perception of traits believed to be shared by humans and animals (e.g., dominance). We tested this in four studies.

In Study 1, we assessed the correspondence between ratings of the same faces presented upright and inverted, finding that inversion reduced the consistency of trustworthiness ratings more than the consistency of dominance ratings. In Study 2, participants categorized the dominance or trustworthiness of upright and inverted faces in a speeded categorization task, showing less ability to identify inverted faces as trustworthy than as dominant. We replicated these results in Study 3, which we pre-registered using the Open Science Framework (see public registration at https://osf.io/zjevk/). Finally, in Study 4, we demonstrated that trustworthiness is considered a more uniquely human trait than dominance (Study 4a), and that face inversion disrupts inferences of trustworthiness and humanness more than dominance (Study 4b). Together, these studies provide evidence that face inversion selectively disrupts humanity-related trait inferences.

4. Study 1

In Study 1, we randomly assigned participants to rate the trustworthiness or dominance of individual male faces. Participants rated each face both upright and inverted. This allowed us to calculate correlations for each participant's ratings of upright and inverted presentations of the same identity, and thus test how much judgments of upright faces corresponded to judgments of inverted faces. Consistent with our hypothesis that configural processing plays a greater role in inferences of human-specific traits, we predicted that the correlation between upright and inverted faces would be stronger for the dominance than trustworthiness ratings.

4.1. Method

4.1.1. Materials

We downloaded all 37 Black and 36 White male faces from the Chicago Face Database (version 1.0; Ma, Correll, & Wittenbrink, 2015) and eliminated the last Black target so that we would have equal numbers of each race. We resized the images to 450 × 316 pixels (72 pixels/in.). We selected male faces for the current studies because we did not wish to introduce additional social group factors that might strongly influence trait ratings, and other recent work looking at consistency across multiple presentations also focused on male faces (Hehman, Flake, et al., 2015; Hehman, Leitner, et al., 2015). Moreover, recent work indicates that dominance may be represented differently in male and female faces (Sutherland, Oldmeadow, & Young, 2016)—an issue to which we return in the General Discussion.

4.1.2. Participants and procedure

We recruited 150 American Mechanical Turk (MTurk) Workers so that we could achieve approximately equal Type-I ($\alpha = 0.05$) and Type-II ($\beta = 0.05$) error rates when assuming a medium effect size (i.e., Cohen's $d = 0.6$) in an independent-samples *t*-test. Half of the participants ($n = 75$) rated each target face on trustworthiness from 1 (*Not trustworthy*) to 7 (*Very trustworthy*), whereas the other half ($n = 75$) rated each face on dominance from 1 (*Not dominant*) to 7 (*Very dominant*). Participants saw each of the 72 targets presented upright in one block and inverted in another block for a total of 144 trials. We randomized the presentation of the faces within these blocks and counterbalanced the order of the blocks. We did not collect any additional measures and report all manipulations and exclusions herein.