



Low-level orientation information for social evaluation in face images

Benjamin Balas¹ · M. Quiridumbay Verdugo¹

Published online: 5 February 2018
© Psychonomic Society, Inc. 2018

Abstract

Observers make a range of social evaluations based on facial appearance, including judgments of trustworthiness, warmth, competence, and other aspects of personality. What visual information do people use to make these judgments? While links have been made between perceived social characteristics and other high-level properties of facial appearance (e.g., attractiveness, masculinity), there has been comparatively little effort to link social evaluations to low-level visual features, like spatial frequency and orientation sub-bands, known to be critically important for face processing. We explored the extent to which different social evaluations depended critically on horizontal orientation energy vs. vertical orientation energy, as is the case for face identification and emotion recognition. We found that while trustworthiness judgments exhibited this bias for horizontal orientations, competence and dominance did not, suggesting that social evaluations may depend on a multi-channel representation of facial appearance at early stages of visual processing.

Keywords Face perception · Social cognition · Visual perception

Introduction

A key feature of human social interaction is social evaluation: What will a stranger's personality may be like upon first meeting, and what kinds of behaviors do we think they may engage in? Observers make social inferences based on multiple cues including postural and gestural information (Curhan & Pentland, 2007), vocal cues (Ambady et al., 2002), and facial appearance (Borkenau et al., 2009). The latter has been the subject of a great deal of recent research, which has established that adult observers can make social face evaluations rapidly (Olivola & Todorov, 2010; Todorov, Pakrashi & Oosterhof, 2009; Willis & Todorov, 2006), and these inferences contribute to decisions about governance (Little et al., 2007), criminal sentencing (Eberhardt et al., 2006), and mental health (Fowler, Lilienfeld & Patrick, 2009). These evaluations may not be veridical, but tend to be reliable. Real-world trustworthiness, for example, is not accurately signaled by facial appearance (Rule et al., 2013), despite observers'

reliability in making trustworthiness estimates. Other traits are well predicted by physiognomy, however: Perceived aggression depends on a facial width-to-height ratio that is linked to real-world aggressive behavior (Carre, McCormick & Hariri, 2011; Carre, McCormick, & Mondloch, 2009; Carre & McCormick, 2008). Regardless of whether particular judgments have ecological validity, social face evaluation is important to understand because observers make these judgments frequently, and use them to make real-world decisions about others.

What visual features contribute to how faces are evaluated socially? Depending on the particular judgment under consideration (competence, trustworthiness, etc.) there are multiple studies that demonstrate specific links between other aspects of facial appearance like masculinity, attractiveness, etc. and personality judgments (Little et al., 2012; Macapagal, Rupp, & Heiman, 2011).

Linking social evaluations to high-level properties of faces is useful to some extent because it provides some interesting insight into the covariation between social and physical face properties. The downside, however, is that this often leads to using one complex aspect of facial appearance to explain another. There are several specific proposals regarding the visual information that supports judgments of gender, facial emotion, etc., but many of these are themselves still a challenge to understand in terms of specific computations that we know

✉ Benjamin Balas
benjamin.balas@ndsu.edu

¹ Psychology Department, North Dakota State University, Dept. 2765, P.O. Box 6050, Fargo, ND 58108, USA

are carried out in the visual system. What is missing is a clear link between observers' abilities to make social judgments from face images, and clearly defined visual features that may (or may not) support those judgments. Here, we propose that linking social face evaluations to low-level visual features is an important way to address this fundamental question.

Early cortical visual processing involves cells in primary visual cortex (or V1) that are selective for spatial frequency and orientation (Tootell, et al., 1998). Individual cells usually respond most to edges and bars tilted at a particular angle, and may also respond optimally to stripes of a particular width (coarse vs. fine). Complex images like faces are thus described in early vision by a population code with multiple channels corresponding to different spatial frequencies and different orientation sub-bands. If we are interested in how low-level information contributes to complex judgments, these channels are a natural place to begin. It is straightforward to create stimuli with information from subsets of these channels selectively included or excluded, and measuring observers' recognition ability subject to such manipulations provides a way to determine the utility of the information in a particular channel. A great deal of work has been done using this approach to characterize different face recognition tasks, revealing that for face identification, both middle spatial frequencies (8–16 cycles per face; Costen, Parker & Craw, 1996; Nasanen, 1999;) and orientation channels near horizontal (Dakin & Watt, 2009; Goffaux & Greenwood, 2016) carry particularly important information.

Our goal in the current study was to extend this approach to determine how distinct social judgments (trustworthiness, competence, and dominance) rely on orientation sub-bands (horizontal vs. vertical orientation energy). We chose these examples of social evaluation because all three of these have been studied widely in previous literature, though we do note that attractiveness (along with trustworthiness and dominance) also emerges in some three-factor solutions that account for variance in social evaluation data (Johnson et al., 2010). We chose to not include attractiveness in the current study because it is in large part a physical aspect of face appearance rather than a truly social evaluation. The current set of social judgments have all been proposed as key factors in face evaluation, and all represent social inferences based on physiognomy. We predicted that like many other recognition tasks, all three of these judgments would depend on horizontal orientation energy more than vertical orientation energy. Furthermore, we predicted that this effect would be more pronounced when judgments were more challenging to make (requiring distinctions between faces that are more similar in perceived personality traits). To test these predictions, we asked naïve participants to complete a 2AFC face evaluation task that required them to determine which of two faces exhibited a particular trait the most (e.g., competence) subject to image filtering that selectively removed orientation

information from horizontal or vertical sub-bands. Briefly, we found that contrary to our hypothesis, different social evaluation tasks rely on orientation sub-bands differently. In terms of low-level visual features, information about personality characteristics is distributed across multiple channels in a way that information about identity is not. Moreover, social vision is not homogenous in terms of the information that observers rely upon. We discuss these outcomes in terms of the neural mechanisms that may support such judgments, and known relationships between high-level face properties that correlate with specific personality judgments, and the low-level feature dependence we have observed here.

Methods

Participants

We recruited 68 adult participants to take part in three tasks (Trustworthiness task, $n=19$; Dominance task, $n=26$, Competence task, $n=23$) from the NDSU Psychology Undergraduate Study Pool. All of these participants reported that they had normal or corrected-to-normal vision, no history of neurological problems, and were naïve to the purpose of the experiment. All participants provided written informed consent prior to participation, and were compensated with course credit for volunteering. Recruitment, consent, and testing procedures were approved by the NDSU IRB.

Stimuli

To ensure adequate control over how trustworthy, dominant, and competent the faces in our experiment appeared, we used stimulus sets previously described in Todorov et al. (2013). Specifically, we chose to use sets of “Maximally Distinct Identities” (available at tlab.princeton.edu/databases) depicting individual faces that vary according to the extent to which they appear to exhibit the target characteristic for each trait judgment. Each set of identities included 25 different individuals, and seven levels of each trait ($\pm 1, 2,$ and 3 standard deviations away from neutral appearance, and neutral appearance itself) from which we chose to exclude the neutral images for each identity. This yielded a total of 150 unique images for each task. All of these images were converted to grayscale prior to the filtering operations described below. This is consistent with multiple prior studies examining these social judgments with face images (see Ballew & Todorov, 2007; Dotsch & Todorov, 2012; Little et al., 2007; Oosterhof & Todorov, 2009; for examples) and simplifies the filtering procedures we applied to create our experimental stimuli. Specifically, using grayscale images makes it easier to match the contrast and mean luminance of our stimulus images across conditions and tasks.

To selectively remove information from specific orientation sub-bands, we applied techniques previously described in Dakin & Watt (2009). Briefly, each image was transformed with a Fast Fourier transform into the frequency domain and subsequently we applied a windowing function to the original power spectrum for each image. This windowing function was defined by a Gaussian envelope centered on the target orientation (0° or 90°) with a standard deviation of 20°). Following the application of this windowing function, the manipulated power spectrum was then used to invert the Fourier Transform, yielding an image with orientation energy outside the windowing function excluded (Fig. 1). We created three versions of each image: (1) An image with horizontal orientation energy only, (2) an image with vertical orientation energy only, and (3) an image with both the horizontal and vertical sub-bands included.

Procedure

In each task, we wished to examine how the available orientation sub-bands affected the ability to judge relative trustworthiness, competence, or dominance as a function of how different test faces appeared to observers on the relevant dimension. To do so, we asked participants to make a series of 2AFC trait judgments in which they indicated which of two simultaneously presented faces exhibited the target trait to the greatest degree.

On each trial, participants were presented with two faces belonging to the same “Identity” in our stimulus set (one to the left of center and one to the right), but varying according to the

level of the task-relevant trait. Specifically, we presented participants with “Easy,” “Moderate,” and “Difficult” trials (Fig. 1) depicting faces drawn from the endpoints of the trait spectrum (± 3 SD for “Easy” trials), the middle of the full range (± 2 SD for “Moderate” trials), or close to the neutral face (± 1 SD for “Difficult” trials). The orientation energy available in the pair of images was also varied such that both faces either retained horizontal orientation energy, vertical orientation energy, or both. Task difficulty and orientation energy both varied pseudo-randomly across trials, and each pair of faces was presented twice with the left/right position of the target face (e.g., the most trustworthy looking image in the pair) balanced across trials. Participants were thus presented with a total of 450 trials. Participants were given unlimited time to view each image pair and make their response, and we did not provide online feedback during the course of the experiment. Experiment trials were presented on $1,024 \times 768$ LCD monitors positioned approximately 40cm from the observer, though neither head position nor eye movements were constrained or measured. At this distance, each image subtended approximately $4\text{--}5^\circ$ of visual angle. All stimulus display and response collection routines were implemented using the Psychtoolbox v3.0 extensions for Matlab (Pelli, 1997; Brainard, 1997).

Results

For each participant, we calculated accuracy as a function of orientation energy (horizontal, vertical, and both orientations)

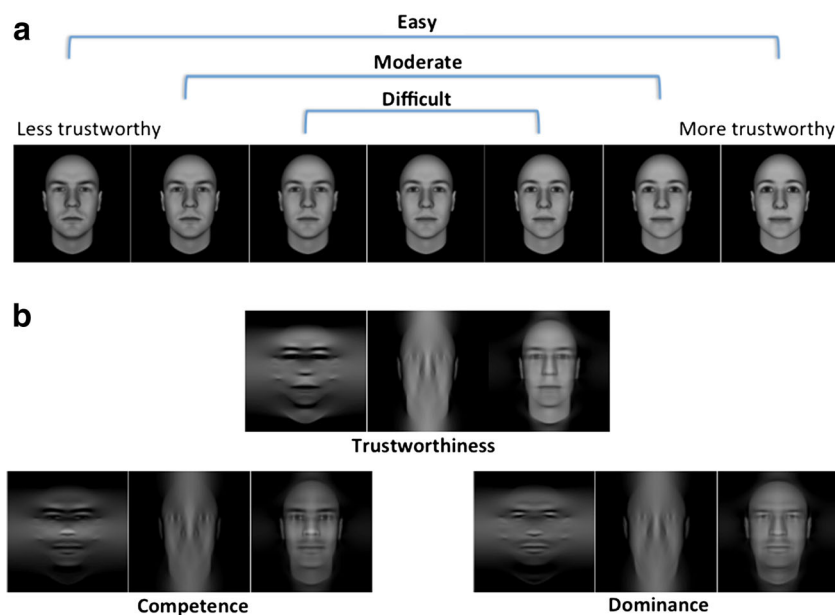


Fig. 1 (a) A single continuum spanning less trustworthy to more trustworthy appearance of the same face identity. The brackets indicate the three levels of difficulty used in all of our tasks. (b) Examples of a

single face from each task filtered to selectively include horizontal orientation energy, vertical orientation energy, or both orientation sub-bands

and difficulty (easy, moderate, difficult). We submitted these values to a mixed-design $3 \times 3 \times 3$ ANOVA with task (trustworthiness, competence, and dominance judgments) as a between-subjects factor and orientation energy and difficulty as within-subjects factor. This omnibus analysis revealed significant main effects of task difficulty ($F(2,128)=65.6$, $p<0.001$, $\eta^2=0.51$), orientation energy ($F(2,128)=22.1$, $p<0.001$, $\eta^2=0.26$), and task ($F(2,64)=19.4$, $p<0.001$, $\eta^2=0.38$). The main effect of task difficulty was the result of significant differences between all three levels of this factor (Bonferroni-corrected post-hoc comparisons all yielded confidence intervals for the difference of pairwise means that did not include zero). The main effect of orientation energy was similarly the result of significant pairwise differences between all three levels of this factor, with performance being lowest in the vertical orientation condition ($M=0.74$, 95 % CI = 0.70–0.77, intermediate in the Horizontal orientation condition ($M=0.80$, 95 % CI = 0.75–0.84) and highest when both orientation sub-bands were available ($M=0.87$, 95 % CI = 0.84–0.90). Finally, the main effect of task was the result of significantly better performance in the Dominance task ($M=0.93$, 95 % CI = 0.88–0.98) relative to both the Competence task ($M=0.74$, 95 % CI = 0.69–0.79) and the Trustworthiness task ($M=0.73$, 95 % CI = 0.67–0.79).

These main effects were qualified by significant two-way interactions between task and task difficulty ($F(4,128)=6.20$, $p<0.001$, $\eta^2=0.16$) and between task and orientation energy ($F(4,128)=4.69$, $p=0.001$, $\eta^2=0.13$). No other interactions reached significance. To examine the nature of these interactions, we chose to conduct separate follow-up ANOVAs for each task separately, to determine how orientation energy and task difficulty affected performance in each case. Specifically, we carried out separate 3×3 repeated-measures ANOVAs for each task, with both orientation energy and task difficulty as within-subjects factors. In Fig. 2, we plot the average accuracy in each task as a function of both stimulus parameters.

Trustworthiness

In our the Trustworthiness task, this analysis revealed main effects of both task difficulty ($F(2,34)=58.2$, $p<0.001$, $\eta^2=0.77$) and orientation energy ($F(2,34)=25.1$, $p<0.001$, $\eta^2=0.60$). The effect of task difficulty was the result of significant pair-wise differences between all three levels of difficulty (Bonferroni-corrected post-hoc comparisons all yielded confidence intervals for the difference of pairwise means that did not include zero). The main effect of orientation energy was the result of significantly lower performance in the vertical condition ($M=0.62$, 95 % CI = 0.57–0.66) relative to both the horizontal condition ($M=0.76$, 95 % CI = 0.70–0.82) and the horizontal+vertical condition ($M=0.81$, 95 % CI = 0.77–0.86). The pairwise comparison between the horizontal condition and the horizontal+vertical condition did not reach significance (95 % CI of the difference between means = -0.01–0.13). Finally, the interaction between task difficulty and orientation energy did not reach significance.

Competence

In the Competence task, we once again found significant main effects of task difficulty ($F(2,44)=5.48$, $p=0.007$, $\eta^2=0.20$) and orientation energy ($F(2,44)=25.1$, $p=0.004$, $\eta^2=0.22$). The interaction between these two factors did not reach significance ($F(4,88)<1$). The main effect of task difficulty was the result of significantly poorer performance in the Difficult condition relative to the Moderate condition (95 % CI of the difference between means = 0.01–0.08, and also relative to the Easy condition (95 % CI of the difference between means = 0.001–0.10). Performance did not significantly differ between the Easy and Moderate conditions (95 % CI of the difference between means = -0.04–0.05). The main effect of orientation energy was the result of significant pairwise differences between performance in the horizontal+vertical condition and both the vertical condition (95 % CI of the difference between means = 0.057–0.30) and also the horizontal condition (95 %

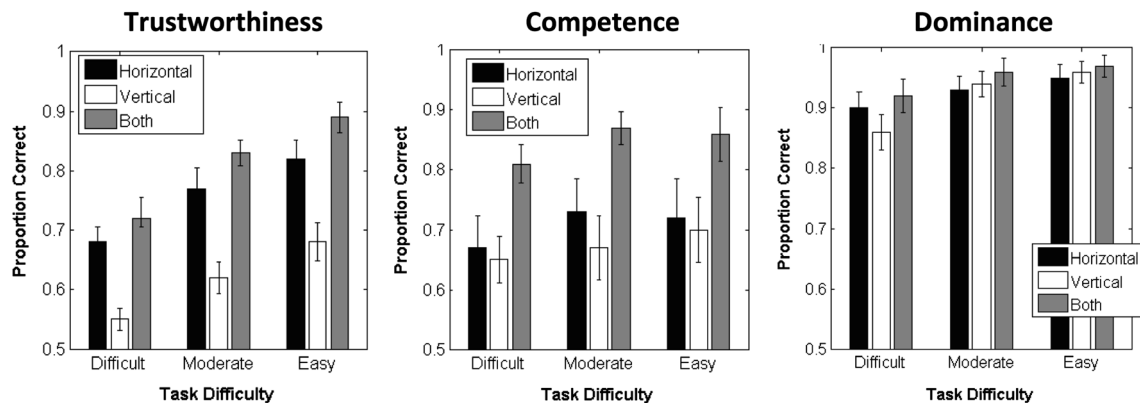


Fig. 2 Average performance in all three tasks as a function of difficulty and orientation condition. Error bars indicate ± 1 standard error of the mean (SEM)

CI of the difference between means = 0.005–0.28). The difference in performance between the horizontal and vertical conditions did not reach significance (95 % CI of the difference between means = -0.19–0.12).

Dominance

Finally, performance in the Dominance task also revealed significant main effects of task difficulty $F(2,50)=30.1$, $p<0.001$, $\eta^2=0.55$) and orientation energy ($F(2,50)=25.1$, $p=0.012$, $\eta^2=0.16$). The main effect of task difficulty was the result of significant pairwise differences in performance between all three levels of this factor (95 % CIs of the difference between means did not include zero in any case). The main result of orientation energy was due to a significant pairwise difference between performance in the horizontal+vertical condition relative to the vertical condition (95 % CI of the difference between means = 0.008–0.053). No other pairwise comparisons reached significance.

These main effects were also qualified by a significant two-way interaction between task difficulty and orientation energy ($F(4,100)=3.90$, $p=0.006$, $\eta^2=0.14$). Post-hoc tests revealed that this interaction was driven by the absence of any differences between orientation conditions at the Easy or Moderate levels, but a significant pairwise difference between the horizontal+vertical condition and the vertical condition at the highest level of difficulty (two-tailed paired-samples *t*-test, $t(25)=3.31$, $p=0.0028$; 95 % CI = 0.022–0.096). We note that compared to the other two tasks, average performance is very high even in the most difficult implementation of the task.

Discussion

These results demonstrate that the visual information that is used for social evaluation in face images depends critically on the specific trait that observers are trying to evaluate. We predicted at the outset that all three judgments would lead to a pattern of behavior consistent with the preferential use of horizontal orientation information for face identification (Dakin & Watt, 2009), emotion recognition (Huynh & Balas, 2014), and other tasks (Goffaux & Dakin, 2010), but found instead that the removal of horizontal or vertical image structure had different consequences for different tasks.

Performance in our trustworthiness judgment yielded results most consistent with our initial hypothesis, demonstrating that observers rely critically on horizontal information: The removal of these features greatly impairs performance, and there is only a minimal cost to presenting observers with only this information. To judge relative trustworthiness from face images, horizontal information appears to be both necessary and sufficient. This outcome can be easily understood via previous results suggesting both that perceived

trustworthiness is affected by subtle variation in emotional expression in neutral faces (Engell, Todorov, & Haxby, 2010; Said, Sebe & Todorov, 2009) and that successful decoding of emotional expressions depends on horizontal orientation energy (Huynh & Balas, 2014; Yu, Chai & Chung, 2011). To the extent that judging trustworthiness amounts to evaluating the valence of facial expressions, it is not surprising to find a common reliance on the same subset of visual features for both judgments.

The data from our competence and dominance judgments are perhaps more compelling, however, precisely because they deviate from what we might expect. In the latter case (dominance) the data are more difficult to interpret because of observers' overall high accuracy in this task, but performance in both tasks differs from our hypothesis. For competence, the removal of either orientation sub-band incurs a consistent cost relative to the control condition where both orientations are present. By contrast, for dominance, neither manipulation has substantial consequences. Though we did observe an effect of filter orientation in the dominance task such that performance in the vertical condition was significantly lower than performance when both orientations were present, the magnitude of that difference was significantly smaller in that task ($M=0.031$, 95 % CI = 0.008–0.053) relative to the trustworthiness task ($M=0.20$, 95 % CI = 0.13–0.27) and the competence task ($M=0.18$, 95 % CI = 0.057–0.030). Both perceived competence and dominance incorporate information from both horizontal and vertical orientation sub-bands. What differs between the two tasks is that for competence that multi-channel representation is fragile (requiring both feature sets) while for dominance it is sturdier (able to rely on either feature set in the absence of the other). Information for dominance may be coded redundantly such that there are common higher-order features available for this task in each sub-band, while the information that is spread across channels for competence judgments is more independent. In terms of the necessity and sufficiency of orientation sub-bands, we can therefore say that for dominance each sub-band is sufficient and not necessary, while for competence, each is necessary and not sufficient.

Characterizing social face evaluation via the low-level features that contribute to distinct tasks supports further exploration of the neural representations that support different judgments. The differential reliance on distinct feature subsets reported here implies that distinct neural loci support these judgments. The common reliance on horizontal orientation energy for trustworthiness judgments and emotion recognition points to a common neural representation supporting observers' performance in both tasks. What about competence and dominance judgments? Both of these seem to require a broadband description of facial appearance such that useful information in both horizontal and vertical sub-bands is available to

participants, with the critical difference in the necessity of each sub-band as a function of task. This may mean that distinct neural computations are used to make these judgments, involving other loci within the extended face network that retain information from a broad range of low-level channels. One possibility is that performance in these tasks may be supported by regions that encode eye gaze, which also depends on vertical orientation energy (Goffaux & Okamoto-Barth, 2013) or that these judgments may depend on the integration of information across a distributed set of neural loci that contribute to face processing. To the extent that both dominance and competence co-vary to some extent with the perceived gender of an individual face (Batres, Re, & Perrett, 2015; Little et al., 2012), this may point to distributed processing of these social judgments: Perceived face gender is signaled by a variety of relatively weak cues including 3D shape (Bruce et al., 1993), 2D pigmentation, and configural cues (Bruce & Langton, 1994), which in turn likely depend on a broad range of computations in the face network.

Another important consideration regarding our results, and the potential for social evaluations to depend on computations across a broad face network, is the extent to which variation in face color, face viewpoint, and other sources of appearance variability may impact the effects we reported here. Our use of grayscale stimuli in the current study leaves open the possibility that the presence of color information may have led to a different pattern of results. Specifically, if there are joint color-orientation features in face images that carry useful information for the social judgments considered here, we may find different outcomes if color information is retained. In this specific case, we feel that our use of grayscale stimuli is well-justified by the multiple prior studies that have used grayscale face images to examine trustworthiness, competence, and dominance (see the Methods section, and also Dzhelyova, Perrett, & Jentsch, 2012; Little et al., 2007; Sussman, Petkova & Todorov, 2013) as well as some evidence suggesting that even when color features (eye color in particular) appear to signal trustworthiness, it is actually shape features predicted by eye color that carry the diagnostic information (Kleisner et al., 2013). However, there is a larger issue regarding generalizability that we think is a critically important direction for future research to pursue: How does the information used for social evaluation change as face variability is allowed to be larger? Choosing to use well-controlled face databases like the images we used here has clear advantages, but also means we accept reduced variability in viewpoint, the lack of hair, the lack of natural face texture (including variation in local color and pigmentation), as well as many other sources of natural variation. While the use of ambient images to validate dimensional models of social face evaluation suggests reasonably good agreement between studies where variability is high and studies where it is more tightly controlled (Sutherland et al., 2013), when we are attempting to

consider the visual features that contribute to these judgments these differences may be more profound. We therefore emphasize that our results must be interpreted within the context of the variability that we did and did not include in our stimulus set, and highlight the need for extended study of this important question regarding how real, noisy images are evaluated in natural settings.

In this context, our results provide a useful description of how biologically-motivated features contribute to complex social judgments, and may lead to predictions about the neural basis of different social evaluations. In both contexts, our data indicate intriguing dissociations between specific visual features and complex social judgments based on facial physiognomy. Further efforts to define how social evaluations are made in terms of features that are grounded in computations known to be carried out in the ventral visual stream may contribute to a more detailed model of how social vision may either be distinct from person recognition, or may co-opt aspects of person recognition based on learned co-variation between social properties and other properties of face images.

References

- Ambady, N., LaPlante, D., Nguyen, T., Rosenthal, R., Chaumeton, N. & Levinson, W. (2002) Surgeons' tone of voice: A clue to malpractice history. *Surgery*, 1325-8.
- Ballew, C.C. & Todorov, A. (2007) Predicting political elections from rapid and unreflective face judgments. *PNAS*, 104, 17948-17953.
- Batres, C., Re, D. & Perrett, D.I. (2015) Influence of perceived height, masculinity, and age on each other and on perceptions of dominance in male faces. *Perception*, 44, 1293-1309.
- Borkenau, P., Brecke, S., Mottig, C. & Paelecke, M. (2009) Extraversion is accurately perceived after a 50-ms exposure to a face. *Journal of Research in Personality*, 43, 703-706.
- Brainard, D. H. (1997) The Psychophysics Toolbox. *Spatial Vision*, 10, 433-436.
- Bruce, V., Burton, A.M., Hanna, E., Healey, P., Mason, O., Coombes, A., Fright, R. & Linney, A. (1993) Sex discrimination: how do we tell the difference between male and female faces? *Perception*, 22, 131-152.
- Bruce, V. & Langton, S. (1994) The use of pigmentation and shading information in recognizing the sex and identities of faces. *Perception*, 23, 803-822.
- Carre, J.M. & McCormick, C.M. (2008) In your face: Facial metrics predict aggressive behaviour in the laboratory and in varsity and professional hockey players. *Proceedings of the Royal Society B: Biological Sciences*, 275, 2651-2656.
- Carre, J.M., McCormick, C.M. & Hariri, A.R. (2011) The social neuroendocrinology of human aggression. *Psychoneuroendocrinology*, 36, 935-944.
- Carre, J.M., McCormick, C.M. & Mondloch, C.J. (2009) Facial structure is a reliable cue of aggressive behavior. *Psychological Science*, 20, 1194-1198.
- Costen, N. P., Parker, D. M., & Craw, I. (1996). Effects of high-pass and low-pass spatial filtering on face identification. *Perception & Psychophysics*, 58, 602–612.

- Curhan, J.R. & Pentland, A. (2007) Thin slices of negotiation: Predicting outcomes from conversational dynamics. *Journal of Applied Psychology*, 92, 802–811.
- Dakin, S., & Watt, R.J. (2009). Biological “bar codes” in human faces. *Journal of Vision*, 9, 1–10.
- Dotsch, R. & Todorov, A. (2012) Reverse correlating social face perception. *Social Psychology and Personality Science*, 3, 562–571.
- Dzhelyova, M., Perrett, D.I. & Jentsch, I. (2012) Temporal dynamics of trustworthiness perception. *Brain Research*, 1435, 81–90.
- Eberhardt, J.L. Davies, P.G., Purdie-Vaughns, V.J. & Johnson, S.L. (2006) Looking deathworthy: Perceived stereotypicality of Black defendants predicts capital-sentencing outcomes. *Psychological Science*, 17, 383–386.
- Engell, A.D., Todorov, A., Haxby, J.V. (2010) Common neural mechanisms for the evaluation of facial trustworthiness and emotional expressions as revealed by behavioral adaptation. *Perception*, 39, 931–941.
- Fowler, K.A., Lilienfeld, S.O. & Patrick, C.J. (2009) Detecting psychopathy from thin slices of behavior. *Psychological Assessment*, 21, 68–78.
- Goffaux, V., & Dakin, S. (2010). Horizontal information drives the behavioral signatures of face processing. *Frontiers in Psychology*, 1, 1–14.
- Goffaux, V. & Greenwood, J.A. (2016) The orientation selectivity of face identification. *Scientific Reports*, 6, 34204.
- Goffaux, V. & Okamoto-Barth, S. (2013). Contribution of cardinal orientations to the “Stare-in-the-crowd” effect. Poster presented at the Vision Sciences Society 2013 Annual Meeting, Naples, Florida.
- Huynh, C. & Balas, B. (2014) Emotion recognition (sometimes) depends on horizontal orientations. *Attention, Perception, & Psychophysics*, 76, 381–392.
- Johnson, S. K., Podratz, K. E., Dipboye, R. L., Gibbons, E. (2010). Physical attractiveness biases in ratings of employment suitability: Tracking down the “beauty is beastly” effect. *The Journal of Social Psychology*, 150, 301–318.
- Kleisner, K., Priplatova, L., Frost, P. & Flegr, J. (2013) Trustworthy-looking face meets brown eyes. *PLoS:ONE*, 8, e53285.
- Little, A.C., Burriss, R.P., Jones, B.C. & Roberts, S.C. (2007) Facial appearance affects voting decisions. *Evolution and Human Behavior*, 28, 18–27.
- Little, A.C., Roberts, S.C., Jones, B.C. & DeBruine, L.M. (2012) The perception of attractiveness and trustworthiness in male faces affects hypothetical voting decisions in wartime and peacetime scenarios. *Quarterly Journal of Experimental Psychology*, 65, 2018–2032.
- Macapagal, K.R., Rupp, H.A. & Heiman, J.R. (2011) Influences of observer sex, facial masculinity, and gender role identification on first impressions of men’s faces. *Journal of Social and Cultural Evolutionary Psychology*, 5, 92–105.
- Nasanen, R. (1999). Spatial frequency bandwidth used in the recognition of facial images. *Vision Research*, 39(23), 3824–3833.
- Olivola, C.Y. & Todorov, A. (2010) Elected in 100 milliseconds: Appearance-based trait inferences and voting. *Journal of Nonverbal Behavior*, 34, 83–110.
- Oosterhof, N.N. & Todorov, A. (2009) Shared perceptual basis of emotional expressions and trustworthiness impressions from faces. *Emotion*, 9, 128–133.
- Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10, 437–442.
- Rule, N.O., Krendl, A.C., Ivcevic, Z. & Ambady, N. (2013) Accuracy and consensus in judgments of trustworthiness: Behavioral and neural correlates. *Journal of Personality and Social Psychology*, 104, 409–426.
- Said, C., Sebe, N., & Todorov, A. (2009). Structural resemblance to emotional expressions predicts evaluation of emotionally neutral faces. *Emotion*, 9, 260–264.
- Sussman, A.B., Petkova, K. & Todorov, A. (2013) Competence ratings in US predict presidential election outcomes in Bulgaria. *Journal of Experimental Social Psychology*, 49, 771–775.
- Sutherland, C. A. M., Oldmeadow, J. A., Santos, I. M., Towler, J., Michael Burt, D., & Young, A. W. (2013). Social inferences from faces: Ambient images generate a three-dimensional model. *Cognition*, 127(1), 105–118.
- Todorov, A., Dotsch, R., Porter, J.M., Oosterhof, N.N. & Falvello, V.B. (2013) Validation of data-driven computational models of social perception of faces. *Emotion*, 13, 724–738.
- Todorov, A., Pakrashi, M. & Oosterhof, N.N (2009) Evaluating faces on trustworthiness after minimal time exposure. *Social Cognition*, 27, 813–833.
- Tootell, R.B.H., Hadjikhani, N.K., Vanduffel, W., Liu, A.K., Mendola, J.D., Sereno, M.I. & Dale, A.I. (1998) Functional analysis of primary visual cortex (V1) in humans. *Proceedings of the National Academy of Sciences USA*, 95, 811–817.
- Willis, J. & Todorov A. (2006) First impressions: Making up your mind after 100 ms exposure to a face. *Psychological Science*, 17, 592–598.
- Yu, D., Chai, A., & Chung, S. T. L. (2011). Orientation information in encoding facial expressions. Poster presented at the Vision Sciences Society 2011 Annual Meeting, Naples, Florida.