Can naive observers detect suicidality in face images? A replication and extension of
Kleiman and Rule (2013)

Hannah I. Pearson and Benjamin Balas

Authors' Note:

Psychology Department, North Dakota State University. Fargo, ND, USA.

Correspondence regarding this article should be addressed to Benjamin Balas,
Department of Psychology, North Dakota State University, Fargo, ND 58102

Contact: Benjamin.balas@ndsu.edu

Abstract

Naive observers make a wide range of social judgments from face images, including estimations of mental health (e.g. depression, schizotypy). Such judgments may inform interactions between care providers and patients, including diagnosis. Kleiman and Rule (2013) demonstrated that naive observers could categorize static images at above chance levels according to suicidality and that the faces of individuals who died by suicide were rated as more impulsive than control faces. We attempted to replicate these results with a larger sample and extended the original report by implementing a 2AFC paradigm that minimizes response bias. We found that suicidality categorization was above-chance using the original testing paradigm, but performance dropped to chance in our 2AFC task. We found no evidence supporting differences in perceived impulsivity across categories in either paradigm. We discuss these outcomes from theoretical and methodological perspectives, and outline new directions for further research investigating social vision in clinical contexts.

Introduction

Human faces support judgments of physical properties like identity, age, and expression, but observers also estimate psychological and behavioral characteristics from face images. These include properties like trustworthiness (Oosterhof and Todorov, 2008), dominance (Rule et al., 2012a; Toscano et al., 2014), aggression (Stillman et al., 2010; Short et al., 2012), and competence (Olivola & Todorov, 2010). The extent to which these judgments reflect real-world behaviors varies substantially (Sussman et al., 2013; Rule et al., 2012b), but whether or not such inferences about personality and behavior are veridical, they are nonetheless reliable within and across observers (Rule et al., 2012b) and can be linked to specific physiognomic features (Oosterhof and Todorov, 2008; Geniole et al., 2014). These aspects of social vision suggest that there are clear biases that affect how observers use facial appearance to make a range of judgments about the people they encounter, some of which may influence critical decisions regarding criminal sentencing (Eberhardt et al., 2006; Flowe, 2012; Wilson & Rule, 2015) and the election of public officials (Todorov et al., 2005; Olivola & Todorov, 2010). Also, the links that have been made between variation along specific dimensions of appearance and real-world behaviors suggest that in some cases there are meaningful correlations between facial appearance and social behavior. For example, the width-to-height ratio of human faces has been demonstrated to be a significant predictor of real-world aggressive behavior (Carre & McCormick, 2008; Carre et al., 2009), suggesting that there may be a common mechanism leading to both phenotypic outcomes.

One domain in which observer biases regarding facial appearance and meaningful links between physiognomy and social evaluation both have important consequences is the perception of mental health. It is evident across multiple studies that observers make judgments from face images that inform their evaluation of the mental health of individuals they see, and in

some cases these judgments may influence diagnosis (Slepian et al., 2014). For example, Friedman et al. (2007) demonstrated that pathological personality traits could be categorized at above-chance levels from 30-second video stimuli. More recent work confirmed that short videos depicting individuals with Borderline Personality Disorder talking about their interests and opinions were rated more negatively than control videos for a range of behaviors including trustworthiness and cooperativity (Hepp et al., 2018). These dynamic stimuli included more than just the face image, but nonetheless support the hypothesis that clinical judgments regarding mental health can be recovered from visual stimuli. With regard to face stimuli specifically, Daros et al., (2016) reported that the faces of individuals with Borderline Personality Disorder could be distinguished from control faces based on subjective ratings of mental health provided by naive participants. The faces of patients with a diagnosis also differed from control faces along other dimensions including anger, disgust, happiness, and emotional stability. This result clearly demonstrates that aspects of mental health can be successfully inferred by untrained observers even when only static face images are available. This phenomenology extends to other disorders, including depression (Scott et al., 2013; Harati et al., 2019) and schizotypy (Ward & Scott, 2018). Facial appearance thus appears to carry some information regarding mental health outcomes that is accessible to untrained observers.

In the current study, we attempted to replicate a particularly important result regarding the relationship between facial appearance and mental health: The ability of naive observers to detect suicidality in static face images. Kleiman and Rule (2013) used a set of images that included photographs of individuals known to have died by suicide and photographs of control individuals who did not. Across multiple experiments, their key results were that naive participants were capable of distinguishing between these two groups of faces and that the faces of individuals who had died by suicide were also rated as more impulsive. Because

suicide represents a major public health concern, the potential for facial physiognomy to contribute to clinicians' ability to identify individuals who may be at risk for suicidal thoughts and behavior is an important contribution. As such, our goals were (1) to replicate the main results of Kleiman and Rule (2013) using a larger sample, and (2) to extend these results using a different testing paradigm. Using the original stimuli, we carried out two experiments that examined observers' ability to distinguish between the face of individuals who died by suicide and control faces. First, we used the same 1AFC categorization and rating tasks described by Kleiman & Rule (2013) to investigate observers' ability to detect suicidality in face images (Exp. 1A) and their tendency to rate these faces as being more impulsive (Exp. 2A). Next, we implemented 2AFC versions of both of these tasks to examine how well observers could distinguish these two groups of faces apart on the basis of perceived suicidality (Exp. 1B) and perceived impulsivity (Exp. 2B). We included these tasks so that we could examine observers' abilities regarding these groups of faces in an experimental context that was nearly free of response bias. Briefly, while we did replicate Kleiman & Rule's original result regarding observers' ability to distinguish the faces of suicidal individuals from control faces in a 1AFC task, we did not observe differences in perceived impulsivity. Further, we did not observe above-chance performance in either of our 2AFC tasks.

Experiment 1

In our first experiment, our goal was two-fold: First, we wished to replicate Kleiman and Rule's main results concerning the categorization of suicidality from individual face images (Experiment 1A). Additionally, we wished to extend this result by examining the discrimination of faces depicting individuals who died by suicide from faces of individuals who had not (Experiment 1B).

In the latter case, we adopted a 2AFC paradigm to minimize the potential impact of response bias on participant judgments.

Experiment 1A - Can naive observers categorize suicidality from individual face images?

Materials and Methods

Stimuli - In all of the experiments reported here we used the original stimuli described by Kleiman & Rule in their original report. This stimulus set consists of 80 total images, 40 photos of individuals who died by suicide (12 women and 28 men) and 40 depicting individuals who did not. These two image subsets are matched for race (87.5% Caucasian, 7.5% East Asian, 2.5% South Asian, 2.5% Latino) and sex (24 women and 56 men). All images were extracted from yearbooks and image quality is generally somewhat low as a result: Some stippling and tonal unevenness is evident in most images but these artefacts are present in both image subsets. All of the images are grayscale and cropped around around the face to exclude any external features including hair, ears, neck, and any other background information (Figure 1). At the time of the original study, all images were cross-checked to make sure that all in the suicide category had actually died by suicide and that images depicting individuals who had not died by suicide were in fact still alive. All images were also pretested during the original task for attractiveness and emotional expression to ensure that these variables were well-matched across both image subsets.

Figure 1. An example image illustrating the general properties of the images in the Kleiman & Rule (2013) stimulus set. To maintain the privacy of the individuals depicted in the stimulus set, we have not included the actual experimental stimuli in these figures, but instead have substituted placeholder images.

Participants

We recruited 45 participants (26 Female) from the North Dakota State University Undergraduate Psychology Study Pool. This sample size was based on a power analysis using the reported mean and standard deviation of A' in Kleiman and Rule's Experiment 1B, and is adequate to detect an effect of the reported size with 95% power. All participants were between the ages of 18-25 years of age and reported normal or corrected-to-normal vision. We obtained informed consent from all participants prior to the beginning of the experimental session including disclosure that participants in the task would be required to look at yearbook images of people who either did or did not die by suicide and categorize images on this basis. Participants were offered an alternative task if they did not feel comfortable participating, but no participants opted out of the main task. All participants received course credit or extra credit for participating in the study (up to the discretion of the professor).

Procedure

Following the procedures outlined in Kleiman & Rule (2013), the 80 images were presented one at a time in a pseudo-randomized order in the middle of the screen. Participants were instructed to indicate after each image whether they believed the photo depicted an individual who had died by suicide or not (pressing “Z” for “Suicide” and “/” for “No.”). To compensate for participants’ presumably low base rate estimates for the incidence of suicide, we informed participants that they would in fact see images of multiple individuals who had in fact died by suicide and should avoid categorizing all faces as “Not suicide.” Participants were permitted to take as much time as they liked per image, but were also encouraged to reach the best decision they could as quickly as possible. All stimulus presentation and response collection routines were implemented using the Psychtoolbox-3 extensions for MATLAB (Brainard, 1997; Kleiner et al., 2007).

Results

All raw data and aggregate files for analyzing our data using JASP (JASP Team, 2018) are available at <https://osf.io/dbkzh/>. To analyze the data, we calculated the hit rate (proportion of faces from the “Suicide” group categorized correctly) and the false alarm rate (proportion of faces from the control group categorized incorrectly) for each participant. Following the analysis outlined in Kleiman & Rule’s original report, we used these values to calculate A' , a non-parametric measure of sensitivity, and B'' , a non-parametric measure of response bias. The formulas we used to convert hit rate and false alarm rate values into A' and B'' are given below:

$$A' = 1/2 + ((h - f) * (1 + h - f)) / 4 * h * (1 - f)$$

$$B'' = (h * (1 - h) - f * (1 - f)) / (h * (1 - h) + f * (1 - f))$$

Both formulas are modified in cases where the hit rate, h , is less than or equal to the false alarm rate, f . These modifications can be found in MacMillian and Creelman (2005) alongside the original formulas that apply to the case where $h < f$.

We continued by carrying out single-sample t-tests to test the hypothesis that sensitivity (A') was larger than 0.5 (the threshold for chance performance) and that bias (B'') differed from 0 in either direction (which would reflect a bias for either response depending on the sign). With regard to sensitivity, we found that A' values ($M=0.524$, $s.d.=0.086$) did significantly differ from chance ($t(44)=1.91$, $p=0.031$, one-tailed t-test, Cohen's $d=0.28$). With regard to response bias, we found that B'' values ($M=0.026$, $s.d. = 0.06$) also differed significantly from the test value of zero ($t(44)=2.86$, $p=0.006$, Cohen's $d = 0.43$). These outcomes suggest that our participants were able to distinguish the faces of individuals who died by suicide from control faces at an above-chance level, and exhibited evidence of a bias favoring "Not Suicide" responses independent of stimulus category.

We chose to complement the original analysis conducted by Kleiman and Rule (2013) with an analysis based on parametric descriptors of sensitivity and response bias. Specifically, we also calculated the values of d' and the response criterion, c , using the hit rate and false alarm rates calculated for each participant. We carried out single-sample t-tests comparing d' values to a test statistic of zero (reflecting chance performance) and comparing c values to zero (reflection no bias for either response type). With regard to sensitivity, this analysis revealed that d' values were significantly larger than chance ($t(44)=1.90$, $p=.032$, Cohen's $d = 0.28$) and that values of c significantly differed from zero ($t(44)=4.15$, $p<0.001$, Cohen's $d = 0.62$). Specifically,

criterion values were significantly positive, reflecting a bias to label faces as “Not Suicide” regardless of image category.

Experiment 1B

Materials and Methods

Stimuli

We used the same stimuli described in Experiment 1A for this 2AFC version of the task. No changes were made to the images save for the differences in presentation during the task described below.

Participants

We recruited 45 participants (24 Female) to participate in Experiment 1B. This sample size was based on a power analysis informed by the reported effects in Kleiman & Rule’s Experiment 1B and was determined to be adequate to detect an effect of the reported size with 95% power. Procedures for participant recruitment and informed consent were identical to those described in Experiment 1A. All participants were between the ages of 18-25 years of age and reported normal or corrected-to-normal vision. Participants received either extra credit or regular course credit for their participation.

Procedure

Unlike the 1AFC judgment we used in Experiment 1A (“Does this face depict someone who died by suicide?”) in this task we used a 2AFC judgment to minimize response bias. We presented participants with an image pair on each trial, one image depicting a face from the “Suicide” group and the other depicting a matched individual from the “Control” group. We asked participants to indicate via keypress which of the two faces depicted someone who had died by suicide (“Z” for the left image and “/” for the right image).

These matched pairs are described in the original Kleiman & Rule report, where the authors note that control images were selected so that each image in the “Suicide” subset was matched to a control image with similar eye color, complexion, and other appearance variables. We presented these image pairs such that the two faces flanked a central fixation point (Figure 2). We presented the 40 unique pairs of images twice each in this task (making 80 total trials) so that the left/right position of the target face was balanced across the entire testing session. As in Experiment 1A, participants were allowed to take as much time as they liked for each judgment, but were encouraged to make the best decision they could as quickly as possible.

Figure 2. Example trial for Experiment 1A including 1 image for each category (“suicide” and “not suicide”). As before, we have substituted the images used in the task with placeholders to protect the privacy of the individuals depicted in the stimulus set.

Results

All raw data and aggregate files for analyzing our data using JASP (JASP Team, 2018) are available at <https://osf.io/dbkzh/>. For each participant, we calculated the proportion of correct trials. Across participants, we observed a mean proportion correct of 0.497 (s.d. = 0.053). To test for potential deviation from normality, we used a Shapiro-Wilk test, which indicated that the values computed across participants did differ significantly from a normal distribution ($W=0.946$, $p=0.35$). Based on this outcome, we tested the hypothesis that the median of the observed values was greater than a test value of 0.5 using a Wilcoxon signed-rank test. This yielded a test statistic of 359.5, and a p-value of 0.67, indicating that participants did not perform at an above-chance level in this task.

Discussion

In Experiment 1A, we replicated the main outcome of Kleiman & Rule’s previous report: Our participants were able to categorize the faces of individuals who died by suicide at an

above-chance level. Our result was also in good quantitative agreement with the prior report (A' was equal to ~ 0.55 in the original study, while our value was ~ 0.53). Though this effect appears to be small, our replication of this effect with a sample that was 150% of the originally reported sample size provides good support for the original effect.

However, despite this replication of Kleiman and Rule's effect in a 1AFC task, we also found that our participants did not perform at an above-chance level in a 2AFC version of the testing paradigm. Specifically, when faces from the "Suicide" subset were presented alongside matched faces from the control subset, our participants could not reliably select which of these faces depicted someone who had died by suicide. How can we resolve this apparent discrepancy between the outcomes of these two tasks? One analytical question we can ask is whether or not participants' decisions from the 1AFC task would be sufficient to succeed at an above-chance level in the 2AFC version. To answer this question using a simple analysis, we calculated the proportion of participants that labeled each image as "Suicide," yielding a score for each image between 0 and 1. Next, for each image pair used in our 2AFC task, we asked how often comparing these values between each target and control image would lead to the correct answer (is the proportion for the "Suicide" item greater than its matched control image?). We determined that this was the case for 22 of our 40 image pairs, which is not significantly better than chance performance via a binomial test. This demonstrates that the different outcomes in these tasks is not truly discrepant, but instead may be the impact of imposing a near-zero bias on participants with stimuli that are difficult to categorize. We continue by examining a second claim from Kleiman & Rule's original report: Do the faces of individuals who have died by suicide look more impulsive?

Experiment 2

In our second experiment, we wished to examine the results reported by Kleiman & Rule (2013) with regard to the perceived impulsivity of the faces of individuals who died by suicide compared to control faces. As in Experiment 1, we carried out both a direct replication of the original testing paradigm with a larger sample and also an extension of the task based on a 2AFC paradigm.

Experiment 2A - Do the faces of individuals who died by suicide look more impulsive?

Materials and Methods

Stimuli

We used the same stimuli described in Experiment 1A for this 2AFC version of the task. No changes were made to the images save for the differences in presentation during the task described below.

Participants

We recruited 45 participants (23 Female) to participate in Experiment 2A. In this experiment, the analysis reported in Kleiman & Rule (2013) was item-based rather than participant-based, so by increasing the sample size we are increasing the precision of the impulsivity estimate per image rather than increasing the power of the subsequent t-test. Still, this sample size represents a 150% increase over the original report in terms of the number of participants. Procedures for participant recruitment and informed consent were identical to those described in Experiment 1A. All participants were between the ages of 18-25 years of age and reported normal or

corrected-to-normal vision. Participants received either extra credit or regular course credit for their participation.

Procedure

The testing procedure for this task was essentially identical to the description provided for Experiment 1A above. Briefly, we presented participants with a single image at the center of the screen and asked them to indicate how impulsive that face looked to them. Participants responded by using the number keys to indicate a Likert scale rating from “1” (Low impulsivity) to 7 (High impulsivity). Participants were encouraged to take as much time as they needed while making the best evaluation they could quickly.

Results

All raw data and aggregate files for analyzing our data using JASP (JASP Team, 2018) are available at <https://osf.io/dbkzh/>. We analyzed the data obtained in this experiment using the same procedures described in Kleiman & Rule’s original report. Specifically, we calculated the average rating per image across participants, and then compared the values obtained from the “Suicide” category to the values obtained from the control category using a paired-samples t-test. Before carrying out this analysis, we confirmed that the data did not deviate from normality according to a Shapiro-Wilk test ($W=0.96$, $p=0.11$). The subsequent t-test revealed no significant difference ($t(39)=1.14$, $p=0.13$, Cohen’s $d=0.18$) between impulsivity ratings assigned to the “Suicide” subset ($M=3.62$, $s.d. = 0.35$) and the ratings assigned to control images ($M=3.53$, $s.d. = 0.45$).

Experiment 2B - Can the faces of individuals who died by suicide be distinguished from controls via perceived impulsivity?

Materials and Methods

Stimuli

Once again, we used the same stimuli described previously for this 2AFC version of the impulsivity task. No changes were made to the images save for the differences in presentation during the task described below.

Participants

We recruited 44 participants (30 Female) to participate in Experiment 2A. Again, this sample size was determined by using the reported estimates from Kleiman and Rule (2012), which indicated that this would be adequate for detecting an effect of the reported size with 95% power. Procedures for participant recruitment and informed consent were identical to those described in Experiment 1A. All participants were between the ages of 18-25 years of age and reported normal or corrected-to-normal vision. Participants received either extra credit or regular course credit for their participation.

Procedure

The procedure was almost identical to Experiment 1B, described above. Once again a 2AFC presentation method was used to show all 40 pairs, with one image from each category, twice.

For each trial, participants were instructed to choose which of the two faces they felt was more impulsive. Their responses were recorded once using keystrokes (“Z” for the left image and “/” for the right image).

Results

All raw data and aggregate files for analyzing our data using JASP (JASP Team, 2018) are available at <https://osf.io/dbkzh/>. As in Experiment 1B, we calculated the proportion of correct trials per participant. Across participants, we observed a mean proportion correct of 0.510 (s.d. = 0.065). To test for potential deviation from normality, we used a Shapiro-Wilk test, which indicated that the values computed across participants did not differ significantly from a normal distribution ($W=0.98$, $p=0.62$). Based on this outcome, we tested the hypothesis that the mean of the observed values was greater than a test value of 0.5 using a one-tailed, single-sample t-test. This yielded a test statistic of $t(43)=0.64$, and a p-value of 0.26, indicating that participants did not perform at an above-chance level in this task.

General Discussion

In our second experiment, we did not replicate Kleiman & Rule’s original finding: We did not find that the faces in our “Suicide” subset were rated as more impulsive, nor did participants reliably label these faces as appearing more impulsive in a 2AFC paradigm where each target face was compared to its matched control face. We suggest that this may mean identifying high-level aspects of physiognomy that are related to this particular mental health outcome remains a challenging and open question.

Considering both of our experiments together, our replication of the key result from the original paper indicates that assessing suicidality is indeed possible in this particular stimulus

set. This is particularly impressive given that the overall quality of the images was low and the images we used in both experiments were cropped to exclude the external contour, which is a critically important visual feature used for the identification of unfamiliar faces (Ellis et al., 1979) and contributes to some social evaluations made from face images (Robinson et al., 2014; Torrance et al., 2014). Moreover, the amount of time that elapsed between the recording of the photographs used in this study and the death of the individuals who died by suicide is variable. This latter point raises an important question regarding the larger enterprise of attempting to use physiognomy to make evaluations of mental health, especially when a particular judgment involves inferences about specific behaviors or emotional states that may be transient: To what extent should we expect that an individual who engages in suicidal ideation should exhibit evidence of that in arbitrary photographs? The use of Ecological Momentary Assessment (EMA) protocols, demonstrates that suicidal ideation fluctuates over a time scale of hours (Kleiman et al., 2017), which strongly suggests that suicidality has important state-like properties. Acute feelings of hopelessness or burdensomeness (Chu et al., 2015) may lead to transient changes in expressive appearance that may not be evident in arbitrary photos. To the extent that suicidality has important state-like properties that affect facial appearance, examining the ability to detect suicidal ideation in single static photographs represents an important proof-of-concept, but may ultimately miss a large portion of the relevant phenomenology. Ideally, examining the ability to detect suicidal ideation using images collected under EMA protocols, or at least using multiple samples collected over a span of days or weeks could help establish when diagnostic features are available to a greater or lesser extent in face and/or body images of at-risk individuals.

These proposals of course pose challenging implementational and ethical questions to address, but again, given the importance of understanding suicidality more completely and

finding ways to help individuals who struggle with suicidal ideation, examining these effects more broadly is warranted. In particular, using dynamic rather than static stimuli may be beneficial. In multiple studies examining “thin-slice” vision in clinical contexts, short video stimuli led to encouraging results demonstrating that various aspects of mental health could be perceived with relatively little visual data. While it is impressive that anything at all can be done with a single static image, understanding the visual features that support this particular judgment (whether we mean specific facial features or high-level appearance variables like impulsivity, etc.) may be a more accessible goal when a larger class of features is available to observers. With regard to suicidality, recent results suggest that an abnormally slow blink rate may be a useful diagnostic criterion (Joiner et al., 2016), providing an important demonstration that dynamic properties of facial appearance may make a substantial contribution to evaluating the mental health of at-risk individuals. We also suggest that applying techniques for identifying critical regions of the face image for suicidality judgments may be a useful means of linking performance in this task to specific features. For example, application of the “Bubbles” technique is of general use for identifying what image regions support a wide range of recognition judgments (Gosselin & Schyns, 2005), and this method has been applied to several aspects of social vision (Robinson et al., 2014). Finally, we also propose that using response classification methods to quantify the perceptual template used by naive observers to make suicidality judgments could be an especially useful tool for understanding how appearance is linked to mental health evaluations (Gosselin & Schyns, 2003; Mangini & Biederman, 2004; Ethier-Majcher et al., 2013). In general, we suggest that continued efforts to understand how this task (and others like it) are accomplished is of particular importance, and will benefit from the adoption of techniques from vision science, social psychology, and machine learning.

Overall, our data largely support the results reported in the Kleiman & Rule's original manuscript. Naive participants were indeed able to perform above chance when categorizing faces according to suicidality. Though their accuracy was low both in the original report and our replication, we suggest that this argues for continued study of how various aspects of mental health can be evaluated from the way faces look. Understanding the mechanisms that support these judgments (whether they are veridical or not!) will have important consequences for clinical practice, visual recognition research, and public health initiatives.

Acknowledgements

We would like to thank Dr. Nicholas Rule for providing the stimuli used in the original report. Also, thanks to Ganesh Padmanabhan and Dan Gu for technical support, and to Dr. Kathryn Gordon for helpful discussions during the design and implementation of this study.

References

- Brainard, D. H. (1997) The Psychophysics Toolbox, *Spatial Vision*, 10,433-436.
- Carre, J.M. & McCormick, C.M. (2008) In your face: facial metrics predict aggressive behavior 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. & Mondloch, C.J. (2009) Facial Structure is a Reliable Cue of Aggressive Behavior. *Psychological Science*, 20, 1194-1198.

Daros, A.R., Ruocco, A.C. & Rule, N.O. (2016) Identifying Mental Disorder from the Faces of Women with Borderline Personality Disorder. *Journal of Nonverbal Behavior*, 40, 255-281.

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.

Ellis HD, Shepherd JW, Davies GM (1979) Identification of familiar and unfamiliar faces from internal and external features: some implications for theories of face recognition. *Perception* 8:431–439.

Ethier-Majcher, C., Joubert, S., & Gosselin, F. (2013). Reverse correlating trustworthy faces in young and older adults. *Frontiers in psychology*, 4, 592.

<https://doi.org/10.3389/fpsyg.2013.00592>

Flowe HD (2012) Do Characteristics of Faces That Convey Trustworthiness and Dominance Underlie Perceptions of Criminality? *PLoS ONE* 7(6): e37253.

<https://doi.org/10.1371/journal.pone.0037253>

Friedman, J.N.W., Oltmanns, T.F. & Turkheimer, E. (2007) Interpersonal perception and personality disorders: Utilization of a thin slice approach. *Journal of Research in Personality*, 41, 667-688.

Geniole, S. N., Molnar, D. S., Carré, J. M., & McCormick, C. M. (2014). The facial width-to-height ratio shares stronger links with judgments of aggression than with judgments of trustworthiness. *Journal of Experimental Psychology: Human Perception and Performance*, 40(4), 1526–1541.

Gosselin, F., & Schyns, P. G. (2005). Bubbles: A User's Guide. In L. Gershkoff-Stowe & D. H. Rakison (Eds.), *Carnegie Mellon Symposia on cognition. Building object categories in developmental time* (p. 91–106). Lawrence Erlbaum Associates Publishers.

Gosselin, F., and Schyns, P. G. (2003). Superstitious perceptions reveal properties of memory representations. *Psychol. Sci.* 14, 505–509.

Harati, S., Crowell, A., Huang, Y., Mayberg, H. & Nemati, S. (2019) Classifying Depression Severity in Recovery from Major Depressive Disorder via Dynamic Facial Features. *IEEE Journal of Biomedical and Health Informatics*, 24, 815-824.

Hepp, J., Storkel, L.M., Kieslich, P.J., Schmahl, C. & Niedtfeld, I. (2018) Negative evaluation of individuals with borderline personality disorder at zero acquaintance. *Behaviour Research and Therapy*. 111, 84-91.

Joiner, T. E., Hom, M. A., Rogers, M. L., Chu, C., Stanley, I. H., Wynn, G. H., & Gutierrez, P. M. (2016). Staring down death: Is abnormally slow blink rate a clinically useful indicator of acute suicide risk? *Crisis: The Journal of Crisis Intervention and Suicide Prevention*, 37(3), 212–217.

Kleiman, E. M., Turner, B. J., Fedor, S., Beale, E. E., Huffman, J. C., & Nock, M. K. (2017). Examination of real-time fluctuations in suicidal ideation and its risk factors: Results from two ecological momentary assessment studies. *Journal of Abnormal Psychology, 126*(6), 726–738.

Kleiman, S. & Rule, N.O. (2013) Detecting Suicidality From Facial Appearance. *Social Psychological and Personality Science, 4*, 453-460.

Kleiner, M., Brainard, D., & Pelli, D. (2007) What's new in Psychtoolbox-3? *Perception, 36*, ECVF Abstract Supplement.

Macmillan, N. A., & Creelman, C. D. (2005). Detection theory: A user's guide (2nd ed.). Lawrence Erlbaum Associates Publishers.

Mangini, M., and Biederman, I. (2004). Making the ineffable explicit: estimating the information employed for face classifications. *Cogn. Sci.* 28, 209–226.

Olivola, C. Y., Todorov, A. T. (2010). Elected in 100 milliseconds: Appearance-based trait inferences and voting. *Journal of Nonverbal Behavior, 34*, 83–110.

Oosterhof, N. N., & Todorov, A. (2008). The functional basis of face evaluation. *Proceedings of the National Academy of Sciences of the United States of America, 105*(32), 11087–11092.
<https://doi.org/10.1073/pnas.0805664105>

Robinson, K., Blais, C., Duncan, J., Forget, H., & Fiset, D. (2014). The dual nature of the human face: there is a little Jekyll and a little Hyde in all of us. *Frontiers in psychology*, 5, 139.

<https://doi.org/10.3389/fpsyg.2014.00139>

Rule, N.O., Adams, R.B., Ambady, N. & Freeman, J.B. (2012a) Perceptions of dominance following glimpses of faces and bodies. *Perception*, 41, 687-706.

Rule, N.O., Krendl, A.C., Ivcevic, Z. & Ambady, N. (2012) Accuracy and Consensus in Judgments of Trustworthiness From Faces: Behavioral and Neural Correlates. *Attitudes and Social Cognition*, 104, 409-426.

Scott, N.J., Kramer, R.S.S., Jones, A.L. & Ward, R. (2013) Facial cues to depressive symptoms and their associated personality attributions. *Psychiatry Research*, 208, 47-53.

Short, L. A., Mondloch, C. J., McCormick, C. M., Carré, J. M., Ma, R., Fu, G., & Lee, K. (2012). Detection of Propensity for Aggression based on Facial Structure Irrespective of Face Race. *Evolution and human behavior : official journal of the Human Behavior and Evolution Society*, 33(2), 121–129. <https://doi.org/10.1016/j.evolhumbehav.2011.07.002>

Slepian, M.L., Bogart, K.R. & Ambady, N. (2014) Thin-Slice Judgments in the Clinical Context. *Annual Review of Clinical Psychology*, 10, 131-153.

Stillman TF, Maner JK, Baumeister RF. A thin slice of violence: Distinguishing violent from nonviolent sex offenders at a glance. *Evolution and Human Behavior*. 2010;31:298–303.

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.

Todorov, A., Mandisodza, A.N., Goren, A. & Hall, C.C. (2005) Inferences of competence from faces predict election outcomes. *Science*, 308, 1623-1626.

Torrance, J.S., Wincenciak, J., Hahn, A.C., DeBruine, L.M., Jones, B.C. (2014) The Relative Contributions of Facial Shape and Surface Information to Perceptions of Attractiveness and Dominance. *PLoS ONE*, 9(10), e104415.

Toscano, H., Schubert, T.W. & Sell, A.N. (2014) Judgments of Dominance from the Face Track Physical Strength. *Evolutionary Psychology*, 12, 1-18.

Ward, R. & Scott, N.J. (2018) Cues to mental health from men's facial appearance. *Journal of Research in Personality*, 75, 26-36.

Wilson, J.P. & Rule, N.O. (2015) Facial trustworthiness predicts extreme criminal-sentencing outcomes. *Psychological Science*, 26, 1325-1331.