



OPEN

Using deep learning to predict ideology from facial photographs: expressions, beauty, and extra-facial information

Stig Hebbelstrup Rye Rasmussen¹, Steven G. Ludeke² & Robert Klemmensen³✉

Deep learning techniques can use public data such as facial photographs to predict sensitive personal information, but little is known about what information contributes to the predictive success of these techniques. This lack of knowledge limits both the public's ability to protect against revealing unintended information as well as the scientific utility of deep learning results. We combine convolutional neural networks, heat maps, facial expression coding, and classification of identifiable features such as masculinity and attractiveness in our study of political ideology in 3323 Danes. Predictive accuracy from the neural network was 61% in each gender. Model-predicted ideology correlated with aspects of both facial expressions (happiness vs neutrality) and morphology (specifically, attractiveness in females). Heat maps highlighted the informativeness of areas both on and off the face, pointing to methodological refinements and the need for future research to better understand the significance of certain facial areas.

What can we learn from a face, and what is it about the face that provides us information? Everyday judgments made by untrained laypeople have some degree of accuracy in predicting a range of psychological characteristics, including personality, intelligence, political ideology, and sexual orientation^{1–4}. In these and other areas, however, deep learning approaches have provided significant predictive improvements over layperson judgments^{5–7}.

The success of deep-learning approaches for predicting psychological characteristics is not only socially important given privacy protection concerns, but scientifically intriguing as well. However, to provide deeper psychological insights and to help the public identify ways to protect sensitive information, the use of deep learning approaches should not stop with successful classification. We should also explore *what* it is about a face that communicates the sensitive information. Are deep learning approaches simply confirming and independently detecting previously-posed facial indicators of a given characteristic, or are they picking up on salient facial differences which were previously unrecognized?

The present study explores the relationship between faces and ideology using a large, publicly-available sample containing both facial photographs and ideological data—namely, 3233 Danish political candidates for local office. We use a wide range of techniques, including convolutional neural networks, heat maps, analyses of facial expressions, and assessments of physical characteristics such as masculinity and beauty. By integrating these approaches we go beyond identifying the degree to which faces connect to ideology in two ways: First we identify specific features of the face that connect to the model's predictions for political ideology. Second, we explore limitations to existing approaches which lead to inflated estimates of the association between faces and psychological characteristics.

Faces and ideology. The study of faces in politics has a long history. Early work found that voters perceive their preferred candidates to be more attractive than opposing candidates⁸. Voter perceptions are themselves politically consequential: For example, faces evaluated as more competent were substantially more successful in an election⁹, and attractiveness provides electoral bonuses as well¹⁰.

This literature has generally posited the relevance of faces to politics to be limited to *subjective impressions* of those faces¹¹. A recent study⁷ strongly refocused the argument by demonstrating that algorithmic classification of faces led to substantial accuracy when predicting ideology. This study's focus, however, was to

¹Department of Political Science, Århus University, Aarhus, Denmark. ²Department of Psychology, University of Southern Denmark, Odense, Denmark. ³Department of Political Science, Lund University, Lund, Sweden. ✉email: robert.klemmensen@svet.lu.se

demonstrate the significant privacy threat posed by the intersection of deep learning techniques and readily-available photographs⁷. This seminal work highlighted the need for further study of the topic, particularly with respect to identifying what features of a photograph contribute to successful algorithmic classification of faces.

Such investigations are in part needed simply to confirm the relevance of *faces* per se. No study has yet employed heat-mapping to identify which segments of facial photographs are informing the classifications provided by the deep learning algorithms when it comes to faces being classified as being rightist or leftist. Because standard cropping procedures for facial photographs (including those employed in prior deep learning facial studies) leave non-facial information in the image, heat-mapping is important to discern whether information beyond the face is informing the classification of faces.

The use of heat-mapping extends beyond demonstrating potential needs for methodological refinements, however, in also highlighting which facial features are the most informative. This facilitates a second area of that needs exploration: *Why* are faces informative for ideology? That is, what information is a neural network detecting when it learns how to correctly identify which faces belong to those endorsing a given ideology? The social importance of the question is clear from a privacy protection perspective: Facial photographs are commonly available to potential employers, and those involved in hiring decisions self-declare a willingness to discriminate based on ideology^{12,13}. Members of the public may thus be aided by recognizing what elements of their photographs could affect their chances of employment. But understanding the facial elements responsible for successful predictions of associations between faces and ideology is also important for the study of ideology. Knowing *that* faces detectably differ based on ideology is undeniably important but identifying the specific features of interest may help to illuminate the origins of ideological differences.

Several plausible candidates exist. Some of these concern facial morphology, i.e., stable physical characteristics of the face. For example, politicians on the right have been found to be more attractive than those on the left¹⁰. Higher right-wing attractiveness scores were also observed among young (male) adults in the general population¹⁰, though not among pundits¹⁴. Masculine facial features may also be relevant: Masculine characteristics such as upper body strength appear to be associated with conservative economic preferences¹⁵, and a preference among conservative voters for masculine-looking males is well established^{16,17}.

Facial expressions may also be relevant. A recent study⁷ reported that computer-coding of facial expressions were associated with self-reported ideology, with those on the left showing more surprise and less disgust. This study⁷ focused on quantifying the degree of threat to privacy, reported only the *degree* to which other facial expressions could independently facilitate the classification of ideology, and not the *direction* of any other associations.

However, the observation of an *actual* association between any given expression or aspect of facial morphology does not entail that this association is being detected by the algorithmic processes used to classify the faces by ideology. No study has yet tested whether model-implied ideology correlates with *any* facial characteristic, whether morphological or expressive. As such, the degree to which these features are informing deep-learning approaches and thus represent a privacy threat needs further consideration.

The present study. We begin by exploring how computational neural networks classify political ideology from a single facial photograph. We then use heat maps to examine how non-facial information contributes to these classifications and to identify facial features of interest for future study. We next examine how a range of characteristics of faces—masculinity, attractiveness, and expressions—connect to model-predicted ideology. Our goal is thus not maximizing accuracy per se—had we explored every model architecture and every set of hyperparameters, it is likely that the accuracy presented here would be increased. But the question of *whether* facial photographs can predict whether someone is leftist or rightist is well answered in the affirmative⁷. We thus instead explore *why* those predictions are possible. The present work is therefore contributing to “interpretable AI”. We add to “interpretable AI” by using theory-guided facial features to seek to identify why deep learning can successfully classify faces by ideology. Thereby we move beyond the many such studies which seek to peek within the “black box” of deep learning by using features of the model itself¹⁸. The present work is thus part of the social data science tradition, in which machine learning methods are combined with theories and empirical results from traditional social science.

Methods

Data and code availability statement. The data has been processed in accordance with Danish and European GDPR regulations according to the University of Aarhus Data Controller Unit. The data is based on publicly available data at the website from the Danish Broadcasting Corporation (DR). Those who wish to access the data should contact DR at <https://dr.custhelp.com/app>. Scripts to redo the analyses will be made available on OSF upon publication. The study was not preregistered. The informed consent of participation and publication of the facial images presented in this article was obtained from the participants.

Sample. *Municipal candidates.* Our primary dataset consists of 5230 facial photographs of political candidates from the 2017 Danish Municipal election. These photos are a publicly-available resource, provided for use in the public sphere by the candidates themselves to the Danish Broadcasting Cooperation (DR). Danish municipal elections take place in a non-polarized setting. Candidates running for office in these elections have not been highly selected through competitive elections in party organizations or by participation in high-stakes elections, and they are thus described as the “last amateurs in politics” by Danish political scientists¹⁹.

Each candidate’s ideology (dichotomously scored as left- or right-wing) was assigned based on the party label under which the candidates ran—see the list of specific parties in Supplementary Materials 1. Members of local parties with less-defined ideologies could not be readily assigned an ideological score and so were excluded

prior to any analyses, leaving 4647 candidates, of which 1442 were female. These images were then manually inspected by an author blind to the candidate's name or party. Additional exclusion criteria applied at this time included: (1) the candidate's face was either not located during pre-processing, not shown in the photo, or not provided in sufficient resolution; (2) the photo was not in color (which would prevent pre-processing all images via gray-scaling in the same way); (3) the candidate was the candidates who did not appear to our rater to be of European ethnic origin, as this small subset (118 individuals) skewed very highly towards left-wing parties, with $2.5\times$ more representation there than among right-wing parties.

After applying these exclusion criteria we had 3288 candidate photos remaining—this is our full primary sample. We then separated the males and females before using R to sample each gender into training, validation and test samples with probabilities of 0.7, 0.15 and 0.15. *Ns* for each sample are provided in Table S1.

Because beards can impair some of the subsequent analyses (such as the detection of facial expressions^{20,21}, we produced a “reduced” version of the male training, validation, and test samples in which those with beards ($N=657$) were dropped. This sample is used for our analyses that involve “identifying salient characteristics” (see below).

Replication sample: Danish parliamentarians. To provide a second test for the accuracy of our algorithmic predictions of ideology based on faces, we obtained an additional sample, namely Danish Parliamentarians. The sample was created following the procedure described for the primary sample, omitting the final two steps. That is, other than separating males and females no division of the sample was undertaken. Instead, we used the entire sample as a test sample. Because we used this small participant pool to evaluate our model's predictive accuracy and not to evaluate facial morphology or expressions, we retained those with beards.

Image preparation. To ensure that only the faces and not the use of a certain background or a certain set of clothes are biasing our results, we used the *dlib* library implemented in Python to locate and crop the face. In addition we used the Python *OpenFace* library²² to align the images such that faces are centered and have the same rotation.

Danish political alignment includes a color element (with left- and right-wing parties denoted by red and blue colors, respectively). To prevent color choices from influencing model results we used *OpenCV* to turn all photos into black and white images. An example image that has been cropped and manipulated as described is provided in Fig. 1. This cropping is all that was performed prior to our first analyses, though as described in the results section, we then found an additional cropping step (described further there) was required.

Measures. Facial expressions. We used the *Face API* from Microsoft Azure's Cognitive Services to identify facial expressions. Facial photographs are assigned a score for how much the face in question resembles each of a list of emotional states—e.g., happiness, sadness, surprise. Descriptive statistics for our full primary sample (Table 1) indicated that faces were overwhelmingly scored as indicating happiness (80%), with neutrality indicated as the second most frequent (19%). The remaining expressions were indicated to be infrequent (ranging of 0.00 to 0.01%). While this balance of facial expressions is not necessarily implausible for photos provided in an election context, the results could potentially reflect imperfect scoring procedures of the *Face API*. An attempt to validate the *API* found it very successful at identifying happy and neutral facial expressions, though it was less successful for others (anger and surprise), with very unsuccessful identification of still others such as disgust and fear²³.

Facial morphology. Masculinity. We follow established practice and measure masculinity in male candidates, using the facial width-height ratio (fWHR; obtained from https://github.com/TiesdeKok/fWHR_calculator), which is illustrated in Fig. 2. Faces high in fWHR are typically perceived as aggressive and dominant²⁴ and as having higher rank and status^{25,26}. High fWHR have been found to predict achievements in both the business²⁷ and athletic domains²⁸.



Figure 1. Example of cropping an image.

	Males				Females			
	Mean	SD	Min	Max	Mean	SD	Min	Max
Anger	0.00	0.01	0.00	0.35	0.00	0.00	0.00	0.03
Contempt	0.01	0.03	0.00	0.45	0.00	0.02	0.00	0.25
Disgust	0.00	0.01	0.00	0.10	0.00	0.01	0.00	0.09
Fear	0.00	0.00	0.00	0.02	0.00	0.00	0.00	0.01
Happiness	0.79	0.32	0.00	1.00	0.92	0.19	0.00	1.00
Neutral	0.19	0.30	0.00	1.00	0.08	0.18	0.00	1.00
Sadness	0.00	0.02	0.00	0.35	0.00	0.01	0.00	0.25
Surprise	0.00	0.00	0.00	0.06	0.00	0.00	0.00	0.01

Table 1. Descriptive statistics for facial expressions. The scores are based on the training samples since this is the largest sample and represents what the model has “learned.” Each face is identified as representing a given expression to a given degree. Thus, a given face could be scored as primarily similar to happiness (0.65), but also similar to a neutral expression (0.35).




Figure 2. Illustration of fWHR.

fWHR scores are obtained by dividing the bizygomatic breadth (the length of the white box) by the distance between the top eyelid and the upper lip (the height of the white box)²⁸. Male masculinity scores in the sample averaged 1.96 [$SD = 0.16$].

Attractiveness. To predict facial attractiveness scores for our sample we use the publicly-available SCUT-FBP5500 dataset (<https://github.com/HCIILAB/SCUT-FBP5500-Database-Release>). This dataset consists of 5500 faces rated by 60 raters using a beauty score ranging from 0–4. (Scoring attractiveness as the average value from multiple raters is commonly employed in political science studies on facial attractiveness¹⁰).

To predict attractiveness scores for our study participants we employ the same CNN model that we used for ideology. The only difference between the network used for predicting ideology and this model is that we are faced with a regression problem instead of a classification problem. Consequently the loss-function is therefore different. We create one model for each gender, using only Caucasians. The model for females had a mean squared error (MSE) of 0.41 the model for males had an MSE of 0.35. Attractiveness scores in our sample averaged 3.35 [$SD = 0.26$] for males and 3.24 [$SD = 0.28$] for females on a scale from 0–4.

Analyses. *Convolutional neural networks.* Deep learning approaches typically require vast amounts of information for effective training. Studies using sample sizes similar to ours often rely on pre-trained networks built on exposure to millions of example images, which researchers then fine-tune to the particular application of the study²⁹. We also employ data augmentation, where we randomly resize and add noise to the existing images, which in essence corresponds to highly increasing the number of training examples at our disposal³⁰. The pre-trained network used in the present study is VGG 16³⁰. We use the training samples (which as noted above represent 70% of our full sample N) to fine-tune the network separately for males and for females. The existing learned parameters within VGG 16 (termed “weights” in deep learning research) are used and refined in the process of trying to predict a new outcome, in our case whether an image represents a left- or right-wing individual. The validation samples (15% of the full sample N for each gender) help prevent overfitting, where a model closely fits the training data but has little ability to predict outside of the training sample. Specifically, we select the set of weights developed within the training sample that show the highest accuracy in the validation sample. These weights are then used to predict ideology within the test sample, much like how regression weights in an ordinary least squares (OLS) regression can be used to calculate individual predictions for a given individual in a sample. Since we already know the ideology of these candidates, we can calculate how accurate the predictions are by comparing the predictions to the actual values.

The deep learning architecture we use is a convolutional neural network (CNN). A CNN uses multiple *local* features in an image instead of using all individual features in an image independently³¹. Simplifying somewhat,




Figure 3. Illustrative heat maps for males and females. *Note:* These pictures represent where the CNN paid attention for females and for males—specifically, the average grad cam results for male and female left-wing candidates. Grad cam images for right-wing candidates can be found in Supplementary materials 1 Fig. 2.

a CNN extracts small pieces of information from the total image and uses these local “snapshots” as features in the neural network instead of inputting all the information in the image as individual features at once.

Consequently, CNNs have two important characteristics²⁹. First, the learned patterns are translation invariant, such that if a particular pattern is learned in an image, the neural network will be able to recognize it in another image even if this feature is not located in the same location in the second image. Second, CNNs learn highly abstract (and hierarchically ordered) features. Thus, we do not expect our CNNs to learn specific features (e.g., a large nose) as predictors of political ideology. Instead, they learn only abstract features that are a combination of a tremendous number of small features—the VGG 16 network we use consists of roughly 140 million parameters. This approach aligns with previous research on facial features in the social sciences, which finds that it is the combination of many features that provide the most useful information¹¹.

Heat maps. CNNs have an important advantage over other deep learning architectures. They can produce visual indications—heat maps—of what they learn from²⁹. Consequently, CNNs are less of a “black box” because they allow us to see which parts of an image a network focuses on when making predictions.

These heat maps have two distinct purposes in the present study. First, they allow us to ensure that the CNNs are attending to the intended elements of the image. Given the present study’s focus on faces, it would be problematic if heat maps showed the CNN was heavily influenced by parts of the image that don’t contain a face, such as the corners of the box that surrounds the face. Second, heat maps provide an informal tool to explore which features of a face are the most informative—for example, are all elements of the face providing important information, or is the CNN particularly attending to one feature or another? This information should facilitate further theory generation.

Results

Heat maps and non-facial information. Figure 3 shows the average heat maps for all male and female candidates in the test sample. These are then superimposed on two images to illustrate *which* areas in the face are most important in terms of making a prediction for left-wing and right-wing candidates.

The preprocessing of images is performed as described above, and illustrated in Fig. 1—i.e., with original photos cropped to a small square with the face in the center, the face rotated, and the image gray scaled. Heat maps for females gave no cause for concern, as classification of females as liberal or conservative by the CNN is overwhelmingly done based on facial features. For males, however, elements beyond the face were substantially informing the classification. In particular, the presence, absence, or characteristics of shirt-collars appeared to be consequential. All results reported in the manuscript therefore use male images which were cropped a second time to exclude all non-facial elements from the photo (see Fig. 4).

Predicting ideology using neural networks. Results on the full primary sample provided accuracies of 61% both for males and for females. These numbers are normed such that using a coin-flip to predict the ideology of the person in the image would be expected to succeed 50% of the time. Our models are therefore better at predicting ideology than a random guess. The accuracy results for males were affected by our exclusion of non-facial information. When our CNN was allowed to use versions of the male photographs (for training, validation, and test) that did not exclude non-facial elements—i.e., the version represented in Fig. 1—the accuracy was higher (65%).

To assess the replicability of our results we used the same CNN weights obtained after training and validation in the primary sample to predict ideology in our replication sample of Danish parliamentarians. For males in this sample we removed background information from the photo as described for the primary sample in order



Figure 4. Illustration of male photo with background removed.




Figure 5. Facial averages for males and females. Note. Facial averages for the candidates with the 20 most extreme scores on predicted ideology for each gender. E.g., the facial average of the 20 male candidates with the most extreme predicted left-wing beliefs are top left.

to include information not related to the face. In this sample we obtained accuracies of 61% for females and 57% for males, which is comparable to accuracies obtained in our primary sample. We now turn to exploring what has enabled the CNN this degree of success in identifying the ideology of participants based only on their face.

Identifying salient characteristics. *Facial averages.* We created facial composites⁵ for individuals identified by our model to be most likely to be left- or right wing. We do this to investigate if facial features recognizable by the human eye are responsible for the predictive accuracy of our model. Figure 5 represents the four resulting facial averages, with the top left face, representing the average of the 20 male candidates our model identified as most likely to be left-wing. Most clearly we see that both male and female right-wing composites appeared happier than their left-wing counterparts. The male composites were not obviously differentiated based on masculinity. The right-wing faces, particularly for females, might be perceived as more attractive.

Morphology. Consistent with the impression from facial composites, masculinity scores were unconnected to model-predicted conservatism [$r=0.02$, 95% CI $[-0.10; 0.04]$]. All reported correlations are Spearman correlations to account for the fact that many of our measures are not normally distributed but quite skewed. Attractiveness was not significantly correlated with model-predicted conservatism in males [$r=0.03$, 95% CI $[-0.033; 0.088]$]. By contrast—and also matching the impression conveyed by facial composites—females attractiveness scores positively correlated with model-predicted conservative ideology [$r=0.11$, 95% CI $[0.042; 0.178]$].

Expressions. Consistent with informal impressions from the facial averages, happy facial expressions were associated with model-predicted ideology. The predicted probability of a male face being right-wing correlated 0.17 (95% CI [0.11, 0.23]) with happiness and -0.17 (95% CI $[-0.23; -0.11]$) with neutrality. The corresponding results for females were 0.15 (95% CI [0.08; 0.22]) and -0.15 (95% CI $[-0.22; 0.08]$).

That the results for happiness and neutrality are nearly exact reversals reflects the fact that other facial expressions were very infrequent in this sample. As noted above, the procedure used to score facial expressions appears to be less accurate beyond happiness and neutrality, and so interpreting results for other facial expressions require caution. However, we did observe that among women contempt was non-trivially correlated with a predicted probability of left-wing ideology ($r = 0.12$, 95% CI [0.05, 0.19]). No other relationship between facial expressions and ideology were statistically significant for either gender.

Discussion

Our results confirmed the threat to privacy posed by deep learning approaches. Using a pre-developed and readily available network that was trained and validated exclusively on publicly available data, we were able to predict the ideology of the pictured person roughly 60% of the time in two samples. We also found that existing practices, in which non-facial information is typically included in the small boxes surrounding a face, could produce inflated accuracies: Heat mapping results showed our CNN was originally using information other than the face to classify males. Had we not deleted this information our accuracy would have been higher.

We also provide the first demonstration that model-predicted ideology connects to independently classifiable features of the face. For females (though not males), high attractiveness scores were found among those the model identified as likely to be conservative. These results are credible given that previous research using human raters has also highlighted a link between attractiveness and conservatism¹⁰. Recall that we use photographs provided by the candidates themselves while campaigning for (low-level) political office. Because attractiveness generally helps electoral success, all candidates are incentivized to provide an attractive photograph¹⁰. Further, the amateur nature of these candidates is a strength because the photos represent candidates for office where selective pressures are rather low. Consequently, we are less (which is not to say zero) concerned that the relationship results from a particularly intensive selection of attractive candidates among those on the right.

Attractiveness was not the only correlate of model-predicted ideology. We also found that expressing happiness is associated with conservatism for both genders. Previous work has found smiling in photographs to be a valid indicator of extraversion³², and while extraversion is not broadly associated with ideology³³, some studies have found that right-wing politicians are more extraverted³⁴; though see³⁵. Self-presentation should also be considered. Politicians on the left and right may have different incentives for smiling—for example, smiling faces have been found to look more attractive³⁶, which is comparatively important for conservative politicians. Future work is needed to explore the extent to which happy faces are indicative of conservatism outside of samples of politicians.

While the use of a politician sample represents a potential constraint on generality, the use of candidates for local office in Denmark, attenuates this constraint. First, because selective pressures are relatively small on this candidates and secondly because prior work on attractiveness and politics reports similar findings across different highly developed political contexts¹⁰.

Future work might explore why our heat maps showed the eyes and mouth regions to be important for the classification of female ideology. While the relevance of these regions to identify things such as whether a face is happy may seem straightforward, happiness correlated with model predicted ideology similarly for males and females, so why a given facial region would be more diagnostic for one gender than another is not immediately obvious.

Received: 21 October 2022; Accepted: 17 March 2023

Published online: 31 March 2023

References

- Penton-Voak, I. S., Pound, N., Little, A. C. & Perrett, D. I. Personality judgments from natural and composite facial images: More evidence for a 'kernel of truth' in social perception. *Soc. Cogn.* **24**, 607–640 (2006).
- Zebrowitz, L. A., Hall, J. A., Murphy, N. A. & Rhodes, G. Looking smart and looking good: Facial cues to intelligence and their origins. *Pers. Soc. Psychol. Bull.* **28**, 238–249 (2002).
- Rule, N. O. & Ambady, N. Democrats and republicans can be differentiated from their faces. *PLoS ONE* **5**, 1–7 (2010).
- Samochowiec, J., Wänke, M. & Fiedler, K. Political ideology at face value. *Soc. Psychol. Personal. Sci.* **1**, 206–213 (2010).
- Kosinski, M. & Wang, W. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images | stanford graduate school of business. *J. Personal. Soc. Psychol. Press* **114**, 246–257 (2017).
- Kachur, A., Osin, E., Davydov, D., Shutilov, K. & Novokshonov, A. Assessing the big five personality traits using real-life static facial images. *Sci. Rep.* **10**, 1–11 (2020).
- Kosinski, M. Facial recognition technology can expose political orientation from naturalistic facial images. *Sci. Rep.* **11**, 1–7 (2021).
- Jahoda, G. Political attitudes and judgments of other people. *J. Abnorm. Soc. Psychol.* **49**, 330–334 (1954).
- Todorov, A., Mandisodza, A. N., Goren, A. & Hall, C. C. Psychology: Inferences of competence from faces predict election outcomes. *Science* **308**, 1623–1626 (2005).
- Berggren, N., Jordahl, H. & Poutvaara, P. The right look: Conservative politicians look better and voters reward it. *J. Public Econ.* **146**, 79–86 (2017).
- Todorov, A. *Face Value*. (Princeton University Press, 2017).
- Inbar, Y. & Lammers, J. Political diversity in social and personality psychology. *Perspect. Psychol. Sci.* **7**, 496–503 (2012).
- Honeycutt, N. & Freberg, L. The liberal and conservative experience across academic disciplines: An extension of inbar and lammers. *Soc. Psychol. Personal. Sci.* **8**, 115–123 (2017).
- Lönqvist, J. E. Just because you look good, doesn't mean you're right. *Personal. Individ. Differ.* **108**, 133–135 (2017).

15. Petersen, M. B. & Laustsen, L. Upper-body strength and political egalitarianism: Twelve conceptual replications. *Polit. Psychol.* **40**, 375–394 (2019).
16. Carpinella, C. M., Hehman, E., Freeman, J. B. & Johnson, K. L. The gendered face of partisan politics: Consequences of facial sex typicality for vote choice. *Polit. Commun.* **33**, 21–38 (2016).
17. Laustsen, L. & Petersen, M. B. Winning faces vary by ideology: How nonverbal source cues influence election and communication success in politics. *Polit. Commun.* **33**, 188–211 (2016).
18. Hu, L., Chen, J., Nair, V. N. & Sudjianto, A. Locally Interpretable models and effects based on supervised partitioning (LIME-SUP). pp. 1–15 (2018) arXiv preprint arXiv:1806.00663.
19. Le Madsen, C. & Kjaer, U. *De sidste Sande Amatører: Kommunalpolitikere og Deres Rolle i det Lokale Demokrati.* (Gyldendal, 2015).
20. Li, Y., Zeng, J., Shan, S. & Chen, X. Occlusion aware facial expression recognition using CNN With attention mechanism. *IEEE Trans. Image Process.* **28**, 2439–2450 (2019).
21. Craig, B. M., Nelson, N. L. & Dixon, B. J. W. Sexual selection, agonistic signaling, and the effect of beards on recognition of men's anger displays. *Psychol. Sci.* **30**, 728–738 (2019).
22. Amos, B., Ludwiczuk, B. & Satyanarayanan, M. *Openface: A general-purpose face recognition library with mobile applications.* *Tech. Rep. CMU-CS-16-118* **16**, 20 (2016).
23. Dehghan, A., Ortiz, E. G., Shu, G. & Masood, S. Z. DAGER: Deep age, gender and emotion recognition using convolutional neural network (2017) arXiv preprint arXiv:1702.04280.
24. Stirrat, M. & Perrett, D. I. Valid facial cues to cooperation and trust: Male facial width and trustworthiness. *Psychol. Sci.* **21**, 349–354 (2010).
25. Hahn, T. *et al.* Facial width-to-height ratio differs by social rank across organizations, countries, and value systems. *PLoS ONE* **12**, e0187957 (2017).
26. Maeng, A. & Aggarwal, P. Facing dominance: Anthropomorphism and the effect of product face ratio on consumer preference. *J. Consum. Res.* **44**, 1104–1122 (2018).
27. Wong, E. M., Ormiston, M. E. & Haselhuhn, M. P. A face only an investor could love: CEOs' facial structure predicts their firms' financial performance. *Psychol. Sci.* **22**, 1478–1483 (2011).
28. Tsujimura, H. & Banissy, M. J. Human face structure correlates with professional baseball performance: Insights from professional Japanese baseball players. *Biol. Lett.* **9**, 20130140 (2013).
29. Carré, J. M., McCormick, C. M. & Mondloch, C. J. Facial structure is a reliable cue of aggressive behavior: Research report. *Psychol. Sci.* **20**, 1194–1198 (2009).
30. Chollet, F. Keras: The python deep learning library. *Astrophysics Source Code Library ascl-1806* (2018).
31. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *3rd International Conf. on Learning Representations, ICLR 2015—Conf. Track Proc.* pp. 1–14 (2015).
32. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning.* (MIT Press, 2016).
33. Naumann, L. P., Vazire, S., Rentfrow, P. J. & Gosling, S. D. Personality judgments based on physical appearance. *Pers. Soc. Psychol. Bull.* **35**, 1661–1671 (2009).
34. Sibley, C. G., Osborne, D. & Duckitt, J. Personality and political orientation: Meta-analysis and test of a threat-constraint model. *J. Res. Pers.* **46**, 664–677 (2012).
35. Caprara, G. V., Barbaranelli, C., Consiglio, C., Picconi, L. & Zimbardo, P. G. Personalities of politicians and voters: Unique and synergistic relationships. *J. Pers. Soc. Psychol.* **84**, 849–856 (2003).
36. Nørgaard, A. S. & Klemmensen, R. The personalities of Danish MPs: Trait- and aspect-level differences. *J. Pers.* **87**, 267–275 (2018).

Author contributions

S.H.R.R. conceived and designed the analysis. S.H.R.R. and R.K. collected the data. S.H.R.R. performed the analysis. S.H.R.R., S.L. and R.K. wrote the paper.

Funding

Open access funding provided by Lund University.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-023-31796-1>.

Correspondence and requests for materials should be addressed to R.K.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023