



OPEN

# Facial recognition technology can expose political orientation from naturalistic facial images

Michal Kosinski

Ubiquitous facial recognition technology can expose individuals' political orientation, as faces of liberals and conservatives consistently differ. A facial recognition algorithm was applied to naturalistic images of 1,085,795 individuals to predict their political orientation by comparing their similarity to faces of liberal and conservative others. Political orientation was correctly classified in 72% of liberal-conservative face pairs, remarkably better than chance (50%), human accuracy (55%), or one afforded by a 100-item personality questionnaire (66%). Accuracy was similar across countries (the U.S., Canada, and the UK), environments (Facebook and dating websites), and when comparing faces across samples. Accuracy remained high (69%) even when controlling for age, gender, and ethnicity. Given the widespread use of facial recognition, our findings have critical implications for the protection of privacy and civil liberties.

There is a growing concern that the widespread use of facial recognition will lead to the dramatic decline of privacy and civil liberties<sup>1</sup>. Ubiquitous CCTV cameras and giant databases of facial images, ranging from public social network profiles to national ID card registers, make it alarmingly easy to identify individuals, as well as track their location and social interactions. Moreover, unlike many other biometric systems, facial recognition can be used without subjects' consent or knowledge.

Pervasive surveillance is not the only risk brought about by facial recognition. Apart from identifying individuals, the algorithms can identify individuals' personal attributes, as some of them are linked with facial appearance. Like humans, facial recognition algorithms can accurately infer gender, age, ethnicity, or emotional state<sup>2,3</sup>. Unfortunately, the list of personal attributes that can be inferred from the face extends well beyond those few obvious examples.

A growing number of studies claim to demonstrate that people can make face-based judgments of honesty<sup>4</sup>, personality<sup>5</sup>, intelligence<sup>6</sup>, sexual orientation<sup>7</sup>, political orientation<sup>8-12</sup>, and violent tendencies<sup>13</sup>. There is an ongoing discussion about the extent to which face-based judgements are enabled by stable facial features (e.g., morphology); transient facial features (e.g., facial expression, makeup, facial hair, or head orientation); or targets' demographic traits that can be easily inferred from their face (e.g., age, gender, and ethnicity)<sup>14</sup>. Moreover, their accuracy is relatively low. For example, when asked to distinguish between two faces—one conservative and one liberal—people are correct about 55% of the time (derived from Cohen's *d* reported in Tskhay and Rule<sup>15</sup>), only slightly above chance (50%). Yet, as humans may be missing or misinterpreting some of the cues, their low accuracy does not necessarily represent the limit of what algorithms could achieve. Algorithms excel at recognizing patterns in huge datasets that no human could ever process<sup>16</sup>, and are increasingly outperforming us in visual tasks ranging from diagnosing skin cancer<sup>17</sup> to facial recognition<sup>18</sup> to face-based judgments of intimate attributes, such as sexual orientation (76% vs. 56%)<sup>7,19</sup>, personality (64% vs. 57%; derived from Pearson's *r*s)<sup>20-22</sup>, and—as shown here—political orientation. (For ease of interpretation and comparisons across studies, across this text, accuracy is expressed as the area under the receiver operating characteristic curve (AUC), an equivalent of the Wilcoxon signed-rank test coefficient and the common language effect size).

We used a sample of 1,085,795 participants from three countries (the U.S., the UK, and Canada; see Table 1) and their self-reported political orientation, age, and gender. Their facial images (one per person) were obtained from their profiles on Facebook or a popular dating website. These self-selected, naturalistic images combine many potential cues to political orientation, ranging from facial expression and self-presentation to facial morphology. The ethnic diversity of our sample (it included over 347,000 non-white participants), the relative universality of the conservative-liberal spectrum<sup>23</sup>, and the generic type of facial images used here increase the likelihood that our findings apply to other countries, cultures, and types of images.

Stanford University, Stanford, CA 94305, USA. email: michalk@stanford.edu


AQ1 AQ2 AQ3 AQ4

AQ5

AQ6

Dataset	Country	Political orientation	Number of participants	% Female	% White	Age	
						Median	IQR
Dating website	U.S.	Conservative	463,367	60%	72%	41	[31,53]
		Liberal	399,403	72%	64%	38	[31,50]
	Canada	Conservative	23,407	52%	73%	38	[30,53]
		Liberal	43,715	67%	69%	37	[32,50]
	UK	Conservative	19,604	52%	78%	39	[31,49]
		Liberal	28,281	58%	73%	36	[30,43]
Facebook	U.S.	Conservative	40,905	62%	73%	27	[24,35]
		Liberal	67,113	66%	55%	27	[24,34]
	Total:		1,085,795	65%	68%	37	[29,50]

**Table 1.** Number of participants and the distribution of political orientation, gender, age, and ethnicity. IQR stands for interquartile range.




**Figure 1.** Procedure used to predict political orientation from facial images. (To protect participants’ privacy, we used a photo of a professional model. Their informed consent for publication was obtained).

As we are aiming to study existing privacy threats, rather than develop new privacy-involving tools, we used an open-source facial-recognition algorithm instead of developing an algorithm specifically aimed at political orientation. The procedure is presented in Fig. 1: To minimize the role of the background and non-facial features, images were tightly cropped around the face and resized to 224 × 224 pixels. VGGFace2<sup>24</sup> was used to convert facial images into *face descriptors*, or 2,048-value-long vectors subsuming their core features. Usually, similarity between face descriptors is used to identify those similar enough to likely represent the face of the same person. Here, to identify individuals’ political orientation their face descriptors are compared with the average face descriptors of liberals versus conservatives. Descriptors were entered into a cross-validated logistic regression model aimed at self-reported political orientation (conservative vs. liberal). Virtually identical results were produced by alternative methods: a deep neural network classifier and a simple ratio between average cosine similarity to liberals and conservatives. See the Methods section for more details.

The results are presented in Fig. 2 (blue bars). The accuracy is expressed as AUC, or a fraction of correct guesses when distinguishing between all possible pairs of faces—one conservative and one liberal. In the largest sample, of 862,770 U.S. dating website users, the cross-validated classification accuracy was 72%, which is much higher than chance (50%) and translates into Cohen’s  $d = 0.83$ , or a *large* effect size. (Sawilowsky<sup>25</sup> suggested the following heuristic for interpreting effect sizes: very small [ $d \geq 0.01$ ], small [ $d \geq 0.2$ ], medium [ $d \geq 0.5$ ], large [ $d \geq 0.8$ ], very large [ $d \geq 1.2$ ], and huge [ $d \geq 2$ ].) Similar accuracy was achieved for dating website users in Canada (71%) and in the UK (70%). The predictability was not limited to the online dating environment: The algorithm’s accuracy reached 73% among U.S. Facebook users. To put the algorithm’s accuracy into perspective, consider that human accuracy in similar tasks is 55%, only slightly above chance ( $SD = 4\%$ ;  $CI_{95\%} = [52\%, 58\%]$ )<sup>15</sup>.

Moreover, as shown in Table 2, the algorithm could successfully predict political orientation across countries and samples. Regression trained on the U.S. dating website users, for example, could distinguish between liberal and conservative dating website users in Canada (68%), the UK (67%), and in the Facebook sample (71%). Overall, the average out-of-sample accuracy was 68%, indicating that there is a significant overlap in the links between facial cues and political orientation across the samples and countries examined here.

Both in real life and in our sample, the classification of political orientation is to some extent enabled by demographic traits clearly displayed on participants’ faces. For example, as evidenced in literature<sup>26</sup> and Table 1, in the U.S., white people, older people, and males are more likely to be conservatives. What would an algorithm’s accuracy be when distinguishing between faces of people of the same age, gender, and ethnicity? To answer this question, classification accuracies were recomputed using only face pairs of the same age, gender, and ethnicity.



**Figure 2.** Accuracy of the facial-recognition algorithm predicting political orientation. All 95% confidence intervals are below 1% and are thus omitted. Humans’ and algorithms’ accuracy reported in other studies is included for context.

	Dataset	Accuracy (AUC) of models trained on:			
		1	2	3	4
1	Dating website (U.S.)	72%	68%	65%	69%
2	Dating website (Canada)	68%	71%	63%	65%
3	Dating website (UK)	67%	65%	70%	66%
4	Facebook (U.S.)	71%	66%	65%	73%


**Table 2.** Classification accuracy across the subsamples (rows) and models trained on each of the samples (columns). All 95% confidence intervals are below 1% and are thus omitted.

The results are presented in Fig. 2 (red bars). The accuracy dropped by only 3.5% on average, reaching 68%, 68%, 65%, and 71% for the U.S., Canadian, and UK dating website users, as well as the U.S. Facebook users, respectively. This indicates that faces contain many more cues to political orientation than just age, gender, and ethnicity.

Another factor affecting classification accuracy is the quality of the political orientation estimates. While the dichotomous representation used here (i.e., conservative vs. liberal) is widely used in the literature, it offers only a crude estimate of the complex interpersonal differences in ideology. Moreover, self-reported political labels suffer from the reference group effect: respondents’ tendency to assess their traits in the context of the salient comparison group. Thus, a self-described “liberal” from conservative Mississippi could well consider themselves “conservative” if they lived in liberal Massachusetts. Had the political orientation estimates been more precise (i.e., had less error), the accuracy of the face-based algorithm could have been higher. Consequently, apart from considering the absolute classification accuracy, it is useful to compare it with one offered by alternative ways of predicting political orientation. Here, we use personality, a psychological construct closely associated with, and often used to approximate, political orientation<sup>27</sup>. Facebook users’ scores on a well-established 100-item-long five-factor model of personality questionnaire<sup>28</sup> were entered into a tenfold cross-validated logistic regression to predict political orientation.

The results presented in Fig. 3 show that the highest predictive power was offered by openness to experience (65%), followed by conscientiousness (54%) and other traits. In agreement with previous studies<sup>27</sup>, liberals were more open to experience and somewhat less conscientious. Combined, five personality factors predicted political orientation with 66% accuracy—significantly less than what was achieved by the face-based classifier in the same sample (73%). In other words, a single facial image reveals more about a person’s political orientation than their responses to a fairly long personality questionnaire, including many items ostensibly related to political orientation (e.g., “I treat all people equally” or “I believe that too much tax money goes to support artists”).

High predictability of political orientation from facial images implies that there are significant differences between the facial images of conservatives and liberals. High out-of-sample accuracy suggests that some of them may be widespread (at least within Western culture). Here, we explore correlations between political orientation



**Figure 3.** Accuracy afforded by transient facial features and personality traits when predicting political orientation in Facebook users. All 95% confidence intervals are below 1% and are thus omitted.

and a range of interpretable facial features including head pose (pitch, roll, and yaw; see Fig. 3); emotional expression (probability of expressing sadness, disgust, anger, surprise, and fear); eyewear (wearing glasses or sunglasses); and facial hair. Those features were extracted from facial images and entered (separately and in sets) into tenfold cross-validated logistic regression to predict political orientation.

The results presented in Fig. 3 are based on Facebook users (similar results were obtained in other samples; see Supplementary Table S1). The highest predictive power was afforded by head orientation (58%), followed by emotional expression (57%). Liberals tended to face the camera more directly, were more likely to express surprise, and less likely to express disgust. Facial hair and eyewear predicted political orientation with minimal accuracy (51–52%). Even when combined, interpretable facial features afforded an accuracy of merely 59%, much lower than one achieved by the facial recognition algorithm in the same sample (73%), indicating that the latter employed many more features than those extracted here. A more detailed picture could be obtained by exploring the links between political orientation and facial features extracted from images taken in a standardized setting while controlling for facial hair, grooming, facial expression, and head orientation.

An algorithm’s ability to predict our personal attributes from facial images could improve human–technology interactions by enabling machines to identify our age or emotional state and adjust their behavior accordingly. Yet, the same algorithms can accurately predict much more sensitive attributes, such as sexual orientation<sup>7</sup>, personality<sup>20</sup> or, as we show here, political orientation. Moreover, while many other digital footprints are revealing of political orientation and other intimate traits<sup>29–34</sup>, one’s face is particularly difficult to hide in both interpersonal interactions and digital records. Facial images can be easily (and covertly) taken by a law enforcement official or obtained from digital or traditional archives, including social networks, dating platforms, photo-sharing websites, and government databases. They are often easily accessible; Facebook and LinkedIn profile pictures, for instance, are public by default and can be accessed by anyone without a person’s consent or knowledge. Thus, the privacy threats posed by facial recognition technology are, in many ways, unprecedented.

Predictability of political orientation from facial images does not necessarily imply that liberals and conservatives have innately different faces. While facial expression or head pose, facial hair, and eyewear were not particularly strongly linked with political orientation in this study, it is possible that a broader range of higher-quality estimates of those and other transient features could fully account for the predictability of political orientation. Yet, from the privacy protection standpoint, the distinction between innate and transient facial features matters relatively little. Consistently changing one’s facial expressions or head orientation would be challenging, even if one knew exactly which of their transient facial features reveal their political orientation. Moreover, the algorithms would likely quickly learn how to extract relevant information from other features—an arms race that humans are unlikely to win.



Some may doubt whether the accuracies reported here are high enough to cause concern. Yet, our estimates unlikely constitute an upper limit of what is possible. Higher accuracy would likely be enabled by using multiple images per person; using images of a higher resolution; training custom neural networks aimed specifically at political orientation; or including non-facial cues such as hairstyle, clothing, headwear, or image background. Moreover, progress in computer vision and artificial intelligence is unlikely to slow down anytime soon. Finally, even modestly accurate predictions can have tremendous impact when applied to large populations in high-stakes contexts, such as elections. For example, even a crude estimate of an audience's psychological traits can drastically boost the efficiency of mass persuasion<sup>35</sup>. We hope that scholars, policymakers, engineers, and citizens will take notice.

## Methods

The study has been reviewed and approved by Stanford University's IRB. All methods were carried out in accordance with relevant guidelines and regulations. The preregistration documents can be found at <https://osf.io/y5wr9>. The author's notes are available at: <https://bit.ly/kosinski1>.

*The dating website sample* was provided by a popular dating website in 2017. It contains profile images uploaded by 977,777 users; their location (country); and self-reported political orientation, gender, and age.

Political orientation was measured using a multiple-choice item. Those who selected "Very conservative" (7%) or "Conservative" (20%) were labeled as conservative; those who selected "Very liberal" (7%) or "Liberal" (16%) were labeled as liberal. Those who selected "Green" (1%), "Libertarian" (2%), "Other" (5%), "Centrist" (17%), or did not know (26%) are not included in this sample. (Those response options are reported in a slightly altered form to protect the identity of the data source).

Given that people prefer partners of similar political orientation<sup>36</sup>, there should be little incentive to misrepresent one's views in the context of a dating website. The validity of this variable is also supported by the high accuracy and high external validity of the political orientation classifier.

*The Facebook sample* included public profile images, age, gender, political orientation, and personality scores volunteered by 108,018 U.S. Facebook users recruited through an online personality questionnaire between 2007 and 2012. Participants were rewarded by the feedback on their scores and provided informed consent for their data to be recorded and used in research.

Participants' personality was measured using the 100-item International Personality Item Pool (IPIP) five-factor model of personality questionnaire<sup>28</sup>, with a five-point Likert-style response scale ranging from "strongly disagree" to "strongly agree." The scales' Cronbach's  $\alpha$  reliability equaled: 0.84, 0.92, 0.93, 0.88, and 0.93 for openness, conscientiousness, extraversion, agreeableness, and neuroticism, respectively. Two items measuring openness were excluded from scoring because they were used to measure participants' political orientation (see below).

Participants' political orientation was established using the "political views" profile field and two items from the IPIP personality questionnaire: "I tend to vote for liberal political candidates" and "I tend to vote for conservative political candidates." To be labeled as conservative, participants needed to self-describe as "republican" (10%), "conservative" (12%), or "very conservative" (2%); *and* disagree (11%) or strongly disagree (9%) with the former IPIP item; *and* agree (10%) or strongly agree (12%) with the latter. To be labeled as liberal, participants needed to self-describe as "democrat" (15%), "liberal" (15%), or "very liberal" (5%); *and* agree (16%) or strongly agree (26%) with the first IPIP item; *and* disagree (20%) or strongly disagree (17%) with the latter. Participants that did not meet those criteria were not included in this sample.

*Facial images* were processed using Face++<sup>37</sup> to detect faces. Images were cropped around the face-box provided by Face++ (red frame on Fig. 1) and resized to 224 × 224 pixels. Images with multiple faces, or a face-box narrower than 70 pixels, are not included in our sample.

Additionally, Face++ was used to identify participants' age, gender, and ethnicity (white, black, East Asian, and Indian); head orientation (pitch, roll, and yaw; see Fig. 3); emotional expression (probability of expressing sadness, disgust, anger, surprise, and fear); and the presence of any kind of glasses or sunglasses. The accuracy of these estimates was good. Predicted and self-reported age correlated at the  $R = 0.65$  level ( $p < 0.001$ ; root mean square error equaled three years). Predicted and self-reported gender matched for 96% of participants. Ethnicity estimated by Face++ and a hypothesis-blind research assistant matched for 88% of 27,023 facial images, a subset of the Facebook sample. Two hypothesis-blind research assistants labeled a subset of 300 images from the Facebook sample with estimates of facial expression and head pose. The correlation between their averaged rankings and Face++ estimates was  $r = 0.72$  on average (see Supplementary Table S2).

*Classification* employed L2-normalized face descriptors derived from facial images using the VGGFace2 model in ResNet-50 architecture, originally trained on an independent sample of over 3.3 million facial images<sup>24</sup>. We tested three approaches to measuring faces' relative similarity to faces of liberal and conservative others:

- *Cosine similarity ratio*: For each face, we took the ratio between its average cosine similarity with liberal faces and between its average cosine similarity with conservative faces.
- *Logistic regression*: Face descriptors were entered into LASSO logistic regression<sup>38</sup> aimed at distinguishing between liberals and conservatives. We used a 30-fold cross-validation so that predictions were made by classifiers that have not seen a given participant before. Parameter  $\alpha$  was set to 1; parameter  $\delta$  was fitted separately within each training set using tenfold cross-validation.

- *Neural Networks*: Face descriptors were entered into a 30-fold cross-validated deep neural network aimed at classifying liberals and conservatives. We tested several network architectures, yet the accuracy did not substantially exceed one offered by two previous, much simpler, approaches.

Given that all three methods yielded similar classification accuracies, we decided to employ LASSO logistic regression. It is computationally efficient and well known among social scientists.

*Classification accuracy* is expressed as AUC. Red bars in Fig. 2 represent the accuracy estimated on the conservative-liberal face pairs of the same age (+/- one year), gender, and ethnicity. We employed Face++ estimates of these traits, as they were available for all faces. Similar accuracy (71%) was achieved when using ethnicity labels produced by a research assistant and self-reported age and gender (ethnicity labels were available for a subset of 27,023 images in the Facebook sample).

*Facial hair classifier* was built using VGGFace2 face descriptors. A hypothesis-blind research assistant labeled 10,000 facial images of males for the presence of facial hair (dichotomous variable) from the Facebook sample. A second hypothesis-blind research assistant labeled a subset of 2,000 of these images: The inter-rater agreement equaled 95%. Facial hair was present on 59% of faces.

These manual labels were used to train LASSO logistic regression<sup>38</sup>, employing face descriptors to estimate the probability of a given face to contain facial hair. Parameter  $\alpha$  was set to 1; parameter  $\delta$  was fitted using tenfold cross-validation. The tenfold cross-validated classification accuracy on the training sample equaled AUC = 96%.

### Data availability

The datasets (excluding actual images) and code used to compute the results are available at <https://osf.io/c58d3/>.

Received: 2 September 2020; Accepted: 7 December 2020

### References

1. Santow, E. Emerging from AI utopia. *Science* **368**, 9–9 (2020).
2. Ranjan, R. *et al.* Deep learning for understanding faces: machines may be just as good, or better, than humans. *IEEE Signal Process. Mag.* **35**, 66–83 (2018).
3. Mollahosseini, A., Hasani, B. & Mahoor, M. H. AffectNet: a database for facial expression, valence, and arousal computing in the wild. *IEEE Trans. Affect. Comput.* **10**, 18–31 (2019).
4. Bond, C. F. Jr., Berry, D. S. & Omar, A. The Kernel of truth in judgments of deceptiveness. *Basic Appl. Soc. Psychol.* **15**, 523–534 (1994).
5. Wolfhchel, K. *et al.* Interpretation of appearance: the effect of facial features on first impressions and personality. *PLoS ONE* **9**, e107721 (2014).
6. Zebrowitz, L. A., Hall, J. A., Murphy, N. A. & Rhodes, G. Looking smart and looking good: facial cues to intelligence and their origins. *Personal. Soc. Psychol. Bull.* **28**, 238–249 (2002).
7. Wang, Y. & Kosinski, M. Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *J. Pers. Soc. Psychol.* **114**, 246–257 (2018).
8. Rule, N. O. & Ambady, N. Democrats and republicans can be differentiated from their faces. *PLoS ONE* **5**, e8733 (2010).
9. Samochowiec, J., Wänke, M. & Fiedler, K. Political ideology at face value. *Soc. Psychol. Personal. Sci.* **1**, 206–213 (2010).
10. Jahoda, G. Political attitudes and judgments of other people. *J. Abnorm. Soc. Psychol.* **49**, 330–334 (1954).
11. Olivola, C. Y., Sussman, A. B., Tsetsos, K., Kang, O. E. & Todorov, A. T. Republicans prefer republican-looking leaders: political facial stereotypes predict candidate electoral success among right-leaning voters. *Soc. Psychol. Personal. Sci.* **3**, 605–613 (2012).
12. Carpinella, C. M. & Johnson, K. L. Appearance-based politics: sex-typed facial cues communicate political party affiliation. *J. Exp. Soc. Psychol.* **49**, 156–160 (2013).
13. Stillman, T. F., Maner, J. K. & Baumeister, R. F. A thin slice of violence: distinguishing violent from nonviolent sex offenders at a glance. *Evol. Hum. Behav.* **31**, 298–303 (2010).
14. Todorov, A., Olivola, C. Y., Dotsch, R. & Mende-Siedlecki, P. Social attributions from faces: determinants, consequences, accuracy, and functional significance. *Annu. Rev. Psychol.* **66**, 519–545 (2015).
15. Tskhay, K. O. & Rule, N. O. Accuracy in categorizing perceptually ambiguous groups. *Personal. Soc. Psychol. Rev.* **17**, 72–86 (2013).
16. Lecun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
17. Esteva, A. *et al.* Dermatologist-level classification of skin cancer with deep neural networks. *Nature* **542**, 115–118 (2017).
18. Lu, C. & Tang, X. Surpassing human-level face verification performance on LFW with GaussianFace. *Proc. 29th AAAI Conf. Artif. Intell.* 3811–3819 (2014).
19. Leuner, J. A Replication Study: Machine Learning Models Are Capable of Predicting Sexual Orientation From Facial Images. *CoRR* abs/1902.1 (2019).
20. 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**, 8487 (2020).
21. Segalin, C. *et al.* What your Facebook Profile Picture Reveals about your Personality. *Proc. 2017 ACM Multimedia Conf. - MM '17* 460–468 (2017). <https://doi.org/10.1145/3123266.3123331>.
22. 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).
23. Tuschman, A. *Our political nature: the evolutionary origins of what divides us* (Prometheus, New York, 2019).
24. Cao, Q., Shen, L., Xie, W., Parkhi, O. M. & Zisserman, A. VGGFace2: A Dataset for Recognising Faces across Pose and Age. in *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)* 67–74 (IEEE, 2018). <https://doi.org/10.1109/FG.2018.00020>.
25. Sawilowsky, S. S. New effect size rules of thumb. *J. Mod. Appl. Stat. Methods* **8**, 597–599 (2009).
26. Pew Research Center. *Wide Gender Gap, Growing Educational Divide in Voters' Party Identification* (2018).
27. 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).
28. Goldberg, L. R. *et al.* The international personality item pool and the future of public-domain personality measures. *J. Res. Pers.* **40**, 84–96 (2006).

29. Kosinski, M., Stillwell, D. J. & Graepel, T. Private traits and attributes are predictable from digital records of human behavior. *Proc. Natl. Acad. Sci.* **110**, 5802–5805 (2013).
30. Youyou, W., Kosinski, M. & Stillwell, D. J. Computer-based personality judgements are more accurate than those made by humans. *Proc. Natl. Acad. Sci.* **112**, 1036–1040 (2015).
31. Park, G. J. *et al.* Automatic personality assessment through social media language. *J. Personal. Soc. Psychol.* **108**, 934–952 (2014).
32. Quercia, D., Kosinski, M., Stillwell, D. & Crowcroft, J. Our Twitter Profiles, Our Selves: Predicting Personality with Twitter. *2011 IEEE Third Int'l Conf. Privacy, Secur. Risk Trust 2011 IEEE Third Int'l Conf. Soc. Comput.* 180–185 (2011). <https://doi.org/10.1109/PASSAT/SocialCom.2011.26>.
33. Bi, B., Shokouhi, M., Kosinski, M. & Graepel, T. Inferring the demographics of search users. *Proc. 22nd Int. Conf. World Wide Web - WWW '13* 131–140 (2013). <https://doi.org/10.1145/2488388.2488401>.
34. Kosinski, M., Bachrach, Y., Kohli, P., Stillwell, D. J. & Graepel, T. Manifestations of user personality in website choice and behaviour on online social networks. *Mach. Learn.* **95**, 357–380 (2014).
35. Matz, S. C., Kosinski, M., Nave, G. & Stillwell, D. J. Psychological targeting as an effective approach to digital mass persuasion. *Proc. Natl. Acad. Sci.* **114**, 12714–12719 (2017).
36. Huber, G. A. & Malhotra, N. Political homophily in social relationships: evidence from online dating behavior. *J. Polit.* **79**, 269–283 (2017).
37. Face++. AI Open Platform. [www.faceplusplus.com](http://www.faceplusplus.com) (2020).
38. Tibshirani, R. Regression shrinkage and selection via the lasso: a retrospective. *Proc. R. Stat. Soc. B Stat. Methodol.* **73**, 273–282 (2011).

### Author contributions

M.K. conducted the research and wrote the manuscript.

### Competing interests

The author declares no competing interests.


### Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-020-79310-1>.

**Correspondence** and requests for materials should be addressed to M.K.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://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) 2020