Gaydar and the Fallacy of Decontextualized Measurement

Andrew Gelman,^a Greggor Mattson,^b Daniel Simpson^c

a) Columbia University; b) Oberlin College; c) University of Toronto

Abstract: Recent media coverage of studies about “gaydar,” the supposed ability to detect another’s sexual orientation through visual cues, reveal problems in which the ideals of scientific precision strip the context from intrinsically social phenomena. This fallacy of objective measurement, as we term it, leads to nonsensical claims based on the predictive accuracy of statistical significance. We interrogate these gaydar studies’ assumption that there is some sort of pure biological measure of perception of sexual orientation. Instead, we argue that the concept of gaydar inherently exists within a social context and that this should be recognized when studying it. We use this case as an example of a more general concern about illusory precision in the measurement of social phenomena and suggest statistical strategies to address common problems.

Keywords: classification; facial appearance; sexual orientation

gaydar, *n.*

Pronunciation: Brit. /'geida:/, U.S. /'gei,dar/

Etymology: Blend of gay adj. and radar *n.*

slang.

An ability, attributed esp. to homosexual people and likened humorously to radar, to identify a (fellow) homosexual person by intuition or by interpreting subtle signals conveyed by appearance or behaviour.

— Oxford English Dictionary (2003)

“Gaydar” colloquially refers to the ability to accurately glean others’ sexual orientation from mere observation. But does gaydar really exist? If so, how does it work?

Our research, published recently in the peer-reviewed journal PLoS ONE, shows that gaydar is indeed real and that its accuracy is driven by sensitivity to individual facial features as well as the spatial relationships among facial features.

— Tabak and Zayas (2012b)

Citation: Gelman, Andrew, Greggor Mattson, and Daniel Simpson. 2018. “Gaydar and the Fallacy of Decontextualized Measurement.” *Sociological Science* 5: 270-280.

Received: March 6, 2018

Accepted: April 1, 2018

Published: May 24, 2018

Editor(s): Jesper Sørensen, Olav Sørensen

DOI: 10.15195/v5.a12

Copyright: © 2018 The Author(s). This open-access article has been published under a Creative Commons Attribution License, which allows unrestricted use, distribution and reproduction, in any form, as long as the original author and source have been credited.

THE science of detecting sexual orientation has experienced something of a renaissance, attracting researchers whose studies garner broad news coverage. The portmanteau of “gay” and “radar” first emerged in print among gay and lesbian comedians in the mid-1990s (for example, DiLallo and Krumholtz 1994), becoming the name of a popular international dating website in 1999 and debuting on U.S. television through the comedies *Will & Grace* and *Futurama* (both in 1999) and

Queer Eye for the Straight Guy (in 2003). Psychologists soon tested for the ability (Shelp 2003), such as when a Harvard undergraduate's 2005 senior thesis garnered coverage in *Psychology Today* with the announcement, "It's true: Some people really do have 'gaydar'" (Lawson 2005). Subsequent studies garnered international media attention, including "Advances in AI Are Used To Spot Signs of Sexuality" (*The Economist* 2017) and "New Study Finds that Your 'Gaydar' Is Terrible" (Feltman 2015).

One problem common to all of these studies, and to the breathless media coverage that followed them (for example, Schramm 2018), was their overestimation of their import outside laboratory or modeled conditions. Stripping away all social context from an inherently social feature consistently produces results that imply the existence of an essential homosexual nature that can be detected and that algorithms can be trained to perform more accurately than humans. We address these conceptual problems below after detailing the errors of statistical extrapolation that underpin them.

Low Population Frequency Makes Gaydar Unreliable

We first review some mathematics that reveal the limitations of any attempt to identify membership in a group that forms only a small subset of the general population.

Let p be the proportion of adults in the United States who identify as gay. For simplicity, let us assume that p equals 4 percent and that everybody is either gay or straight. (It would not be difficult to extend the model to allow for intermediate or other characterizations.) Now suppose that you have a gaydar that emits a continuous signal when you study a person, and the reflection of that signal contains a signature that informs you that the person is gay. Or, to put it more prosaically, you process some number of attributes conveyed by a person that enable a probabilistic classification, assuming you have been accurately trained and can correctly adjust for changes in base rate.

To simplify, suppose you are required to compress your continuous gaydar measurement into a binary go/no-go decision using some threshold that will trade off between false positives and negatives. Let α be the probability that a gay person is correctly classified as gay, and let β be the probability that a straight person is correctly classified as straight; thus with perfect gaydar, $\alpha = \beta = 1$.

Suppose you classify someone as gay; then the probability that he or she actually is gay is $\Pr(\text{gay given they have been classified as gay}) = \frac{p\alpha}{p\alpha + (1-p)(1-\beta)}$. It is a well-known result from conditional probability that when p is low, the misclassification rate is high. For example, if $p = 0.04$ and $\alpha = \beta = 0.9$ (which would be an impressive accuracy in many settings), then $\Pr(\text{gay} \mid \text{classified as gay}) = \frac{0.04 \cdot 0.9}{0.04 \cdot 0.9 + 0.96 \cdot 0.1} = 0.27$; thus, you would actually be wrong nearly three-quarters of the time. Under this system, you would be classifying 13.2 percent of people as gay, which is plausible given that people generally overestimate the proportions of rare events (see Hemenway 1997) and groups, including immigrants, ethnic minorities, and gays (see Sides and Citrin 2007; Srivastava 2011; Newport 2015).

From these considerations, we can also work out how accurate a classifier needs to be in straight and gay populations to achieve any specified classification accuracy of gay people when applied to a population in which the proportion of homosexuals is p . If we want the chance of someone who is classified as gay actually being gay to be greater than 50 percent, then we need $\beta > 1 - \frac{p\alpha}{1-p}$ to hold. When p is small, this means that β , the probability of classifying a straight person correctly, will have to be very close to 1. For example, if 4 percent of the population is gay, then even if we could identify every single gay person perfectly, we would still need to classify around 96 percent of straight people correctly to achieve the required accuracy. The intuition under this result is that because there are so many more straight people than there are gay people, there is an enormous penalty if they are incorrectly classified. This error underlies two recent gaydar studies that also suffered from conceptualizing homosexuality, and the signals sent about it, as asocial phenomena. The point here is not that partial classification is impossible but rather that it will be challenging to have a high probability of correct prediction in any particular case.

Two Gaydar Experiments

In a study that received extensive media attention, Tabak and Zayas (2012a) assessed the abilities of 24 college-student volunteers at identifying the sexual orientations of 400 self-identified gay or straight people based on photographs of faces that excluded images of individuals with eyewear, jewelry, scars, or other “facial alterations.”¹ Half the targets in the study were gay and half were straight, and the students correctly identified sexual orientation 60 percent of the time, which was statistically significantly better than the 50 percent that would be expected by pure guessing.

Tabak and Zayas write that their research “was the first attempt to determine the roles that featural and configural face processing play in snap judgments of sexual orientation from faces,” and it indeed seems to provide a clue about such visual manifestations. But we disagree with their claim that they have shown that “configural face processing significantly contributes to perception of sexual orientation.”

To understand our disagreement, consider several aspects of gaydar as we understand it and which are consistent with the dictionary definition given at the start of this article. Gaydar occurs in social contexts, with information including voice, dress, posture, and even topics of conversation or the places and contexts in which they occur; the Tabak and Zayas study removes all such cues. Gaydar is relevant in settings where gays are a small fraction of the population; in contrast, gay people represented 50 percent of the photos in the experiment under discussion. Indeed, Tabak and Zayas act as if they have removed all social cues, leaving behind nothing but the asocial and objective face, as if photographs are taken in laboratory conditions or as if smiles or eye expressions are not themselves “facial alterations.” Gaydar under these laboratory conditions has been transmuted from the in-group task of identifying fellow members of a rare subgroup in social interactions to a mechanical binary classification task of asocial faces conceived as neutral.

More recently, Wang and Kosinski (2018) performed a similar exercise, this time using a machine learning algorithm to identify faces as gay or straight using 35,326 images scraped from an unnamed dating website. Once again, approximately half of these images are of people who were using the website to find members of the same sex. Their algorithm was able to classify sexual orientation correctly in 70 to 80 percent of a subset of this data set that was held aside while building the classifier. The classifier also worked, although less well, on a set of around 900 faces of white men from Facebook who both identified as looking for a male partner and liked at least two Facebook pages such as “I love being gay” and “Gay and fabulous,” and in other experiments, they also compared to human judges and to computer classification using landmarks relating to facial morphology.

Here we focus on the classification of the images from the dating website. It can be interesting to see what happens to show up in the data (“Lesbians wear baseball caps!” “Gay men have less facial hair!”), but we question their extrapolation from their experimental conditions into speculations regarding inherent associations of homosexuality, for example:

“It is unclear whether gay men were less likely to wear a beard because of nature (sparser facial hair) or nurture (fashion). If it is, in fact, fashion (nurture), to what extent is such a norm driven by the tendency of gay men to have sparser facial hair (nature)? Alternatively, could sparser facial hair (nature) stem from potential differences in diet, lifestyle, or environment (nurture)?” (Wang and Kosinski 2008:254)

Similarly, the authors suggest that the correlation between facial brightness and probability of being gay could be evidence that straight men have higher levels of testosterone. An alternative reading of that result is that gay men are more likely to postprocess their dating profile pictures using a variety of readily available tools and filters, part of the impression management strategies that are rife in online profile photographs (Toma and Hancock 2010; Zytka, Grandhi, and Jones 2014). As Cohen (2017) and Mattson (2017) note, the speculation that Wang and Kosinski engage in is essentially disconnected from the data analysis that is being used as its justification.

Wang and Kosinski report that their goal was to “advance our understanding of the origins of sexual orientation and the limits of human perception,” but their data provide no evidence regarding the former. As for the latter, a challenge is that the people in the photographs may well be intentionally using their facial expressions to communicate sexual orientation, given that people classified in one of the studies were selected based on Facebook “likes” on such pages as “Gay Times Magazine” and “Manhunt” as evidence for actual sexual orientation.

Neither Wang and Kosinski nor Tabak and Zayas produced results that suggest that more than half of the people flagged by their respective laboratory gaydars would be gay if these gaydars had been applied to the general population, and that is consistent with results from Olivola and Todorov (2010) on the error rates of sexual orientation judgments applied in realistic social settings. In stating these limitations, we do not intend to reject the empirical results of Tabak and Zayas (2012a) and Wang and Kosinski (2018). Their classification results are interesting and may well point to new insights regarding variation in facial expression.

Sampling and Social Context

Both the studies under discussion here measure the perception of sexual orientation in a context-free way. Decontextualization—bringing a phenomenon “into the lab” for careful study—is a characteristic step of scientific measurement, but it can cause problems in fields such as ecology and social science, where context is all. Reductionism—breaking a complex phenomenon into simpler parts to enable understanding—is a necessary part of the scientific enterprise, but bracketing the social for an inherently social phenomenon causes its evaporation, not its reduction. In particular, we have three concerns with these laboratory studies of gaydar.

The easiest concern to state is representativeness. A low frequency of facial hair among openly gay men who post to a particular dating website to find other gay men or a finding that “sexual orientation is inferred more easily from women’s vs. men’s faces” (Tabak and Zayas 2012a) may well be telling us more about the samples than about the general population they are presumed to represent. These are what Magnet (2011) calls the “demographic failures” baked into biometric technologies by reductive and biased samples; after an exhaustive review of their failures, she concludes that “human bodies are not biometrifiable.” Given that no census or representative sample exists of images of gay people (or, for that matter, straight people), any statistical analysis will always have to deal with the extrapolation problem, and we recommend using some sort of multilevel model that explicitly allows for variation among and within different subgroups of each population. After all, there is great variation among even heterosexuals, some of whose “hybrid” expressions—such as metrosexual men or working-class Midwestern women—appear to others to be homosexual when they are not (Kayzak 2012; Bridges and Pascoe 2014; Hall 2015).

Our second concern is the way in which gaydar, which was originally framed as an aspect of communication *within* the gay community, has been redefined in the news media as a skill that can be deployed by the general (thus, mostly straight) population. One can distinguish between “active” gaydar (in which members of a subgroup are sending coded messages to each other—what Shelp [2003] calls adaptive gaydar) and “passive” gaydar, in which outsiders catch some of these signals even though they are not the intended recipients. *In either case*, we suspect that many if not most of the distinctive and noticeable characteristics of the subpopulation are the result of active choices by members of that group, not (as assumed in the two articles under discussion) essential attributes derived from the false binary of “nature” or “nurture.”

By taking gaydar into the lab, these research teams have taken the creative adaptation of an oppressed community of atomized members and turned gaydar into an essentialist story of “gender atypicality,” a topic that is related to, but distinctly different from, sexual orientation (for example, Newton 1984; Valdes 1995; Fausto-Sterling 2000). This new story has moved gay people from the protagonists of the story to research objects about whom there is the potential for construction of just-so stories based on gender stereotypes. Again, a certain amount of reduction and objectification is necessary in scientific research, but researchers must be aware of what is lost in these steps in generalization to real-world settings.

Our third concern is the use of gender stereotypes in the deduction of homosexuality. Researchers have long distinguished between same-sex identity, same-sex desires, and same-sex behavior, distinct phenomena that “are imperfectly correlated and inconsistently predictive of each other” (Savin-Williams 2006; see also Laumann et al. 1994). The relationship between gender identity and gender expression is similarly fraught in relation to sexual orientation. Researchers who attend to such distinctions have not found a reliable gaydar among research subjects, such as the finding that “the stereotypic association of feminine looking men as homosexual may confound judgments of sexual orientation” (Valentova et al. 2014). Other research is definitive, finding that “stereotypes casting gays and lesbians as gender ‘inverts,’ in cultural circulation for a century and a half, lead perceivers to use gendered facial cues to infer sexual orientation” (Freeman et al. 2010; see also Rieger et al. 2010). In other words, people can readily detect gender atypicality, but its relationship to homosexuality—and heterosexuality—is unclear. Given the consistent evidence that same-sex identity is organized differently among men and women, and also that bisexuality may be the most common expression of homosexuality among women (Bailey et al. 2016), it is problematic on gender grounds to reduce sexuality orientation to a binary variable, a reduction that is understandable for the simplicity of modeling but should be carefully noted when discussing these studies.

Gaydar is itself a mode of communication that is contingent on social structures and expectations. Colzato et al. (2010) find that subjects who are told that gaydar exists will have it; they also have found that homosexual subjects attend more closely to patterns and details than heterosexuals, which they presume “increases the likelihood to detect perceptual cues indicative of orientation, which again facilitates finding like-minded, social peers, and potential friends and sex mates.” Their research is consistent with the idea of gaydar as a form of communication and not some intrinsic radiation from flaming homosexuals.

The Fallacy of Decontextualized Measurement

The reporting and interpretation of the gaydar experiments suffered from three problems discussed below, which are common in social science research. We can identify all these problems with what might be called *the fallacy of decontextualized measurement*, the idea that science proceeds by crisp distinctions modeled after asocial phenomena, such as unambiguous medical diagnoses (the presence or absence of streptococcus or the color change of a litmus paper). Seeking an on–off decision, normalizing a base rate to 50 percent, and, most problematically, stripping a phenomenon of its social context all give the feel of scientific objectivity while creating serious problems for generalizing findings to the world outside the lab or algorithm.

The first problem we noticed in these studies is the reduction of an (implied) continuous scale to a binary choice. Some people appear clearly gay, others emit some gay signals, and others appear completely straight. Gaydar is on a sliding scale and depends on context; again, the traditional goal is to identify gay people who might be signaling their sexual identities to the in-group while staying hidden from the general population, which is quite a bit different from a study such as

that of Wang and Kosinski (2018), which used participants on a dating site who both self-identified as gay and wanted other people to know this. Some of the problems can be reduced by at least adding a third category to the classification that accounts for cases in which the classifier (human or machine) is unsure. A guiding principle when trying to split a population into two groups is that you always need a third group to account for those individuals who are within a margin of error (Gelman and Park 2009). Applying this “rule of thirds” would not fix the problem of reducing a continuous scale to a binary choice, but it does allow the researchers to better assess and communicate the uncertainty in their method.

Second, these laboratory experiments have a much different base rate than in real-world settings, a point also made by Plöderl (2014) and Cox et al. (2017). It is well known that judgments of uncertainty are contingent on base rate (Kahneman and Tversky 1974), and this is particularly relevant for a concept such as gaydar, which arises in a setting in which the challenge is identifying a small minority in a large population. This effect is obvious in the Tabak and Zayas study—in which orientation was assessed by humans—and it could be partly mitigated by repeating the experiment with only 4 percent of the photographs portraying a homosexual. It could also be instructive to repeat the experiment by using a set of photographs that are only of heterosexuals; if these studies are purely classificatory exercises, then the ability to spot heterosexuals would be as interesting as any other algorithm.

Adapting the Wang and Kosinski experiment to a low base rate presents a slightly different prospect. There is nothing wrong with using a 50/50 sample in the first step of their procedure, which extracts around 4,000 facial features that are used to distinguish between the two cases. As these researchers note, these features can be converted into a classifier by adjusting the base rate. Although these studies may have not intended to study categorization in real life but rather general perceptual capabilities (for example, Bruno, Lyons, and Brewer 2014), such careful discussion is often missing in study framings and discussions by researchers in the news media.

Third, the researchers took a rich real-world phenomenon and abstracted it so much that they removed all its social content (or believe they have; eyes are maintained, as findings that lesbians wear less eye makeup than straight women attest). Gaydar has traditionally existed within a particular social context: a world in which gays are an invisible minority, existing in plain sight and seeking to be inconspicuous to the general population while communicating with others of their subgroup using various coordination practices (for example, Minnelli et al. 1989). Face shapes may tell us something about gender atypicality, and we see value in studies of human and machine classifications of faces, but interpretations of these empirical findings will be made in the context of social norms and the ways humans interact.

Discussion

In recent years psychologists have studied correlates of sexual orientation in a variety of ways that have attracted various media coverage (for example, France 2007; Saletan 2011). The place of homosexuality in our culture has changed dramatically in recent decades, to the extent that concepts such as gaydar are changing their

meaning and probably also their practice. This suggests that the signaling of the past is different from that of today: compare the cagey flamboyance of male figure skater Johnny Weir in the 2008 Winter Olympics to the unabashed “faggy magic” of Adam Rippon in the games one decade later (Moskowitz 2018).

As noted above, we are not claiming that experiments such as those of Tabak and Zayas (2012) and Wang and Kosinski (2018) are useless. If various aspects of the shapes of faces are (weakly) correlated with sexual orientation, even through the intervening variable of gender atypicality, and if untrained volunteers or computer programs can classify such patterns with high accuracy, this is possibly interesting. Some insight might be gained by performing a set of studies comparing other groups, each time using people or computer programs to classify people chosen from two different samples (for example, college graduates and noncollege graduates, English people and French people, driver’s license photos in state X and driver’s license photos in state Y, students from college A and students from college B, baseball players and football players, or people on straight dating site U and people on straight dating site V). From a statistical standpoint, it could be useful to think of gay versus straight as one of many possible classifications, and researchers could then explore the ways in which classification accuracy varies by group division and data source. Reductionism is a characteristic and useful tool of science when used to understand a complex phenomenon by breaking it down into simple parts. In these cases, however, too much has been stripped away from social reality to make any general conclusions about differences between gay and straight people.

The word “gaydar” never appeared in either of the two articles under discussion (except in some of their references), but the expression was all over the place in media reports as well as in critical discussions, such as this one. This illustrates one of the challenges of research in this area: scientific research does not take place in a vacuum. It is not paradoxical that two articles that never use the word “gaydar” can be all about gaydar if this is how they are received. We thus emphasize that we are exploring these research articles and their scientific and popular receptions all at the same time. If popular conceptions (and misconceptions) of gaydar are warping the reception of serious scientific studies, this is something researchers must confront.

The point of this article is not to pick on a small area of psychology research that happened to catch the fancy of the press or even to criticize larger trends of sensationalism in science and the news media. Rather, we seek to draw attention to the general problem (one that is all too frequent in this era of genetics, machine learning algorithms, and MRI studies) in which the ideals of scientific precision end up stripping all context from a social phenomenon, leading to nonsensical claims based on predictive accuracy or statistical significance. Statistical and machine learning methods can be very useful in social science, but they are hard to use well and can be easily misinterpreted. These algorithms are only able to do one thing: they interpolate the training data. It is difficult to bracket social stereotypes, which often form the taken-for-granted common sense, leading to Ceglowski’s (2016) caution that “machine learning is like money laundering for bias” when artificial worlds are substituted for the social world. Indeed, many of our critiques have been leveled by psychologists and artificial intelligence researchers at their colleagues

(Cox et al. 2016, 2017; Fasoli and Hegarty 2017). Restoring context may come from research partnerships, perhaps by having studies reviewed by experts in the social phenomena and not merely the methods or even coauthoring with them. A social interaction cannot always be measured in a test tube or even in a psych lab. In the case of the research under discussion here, the steps taken ostensibly to ensure objectivity obliterated much of the objects of research.

Notes

- 1 "To minimize the prospect that non-face cues would influence judgments, photographs of men or women with facial alterations or adornments (for example, scars, eyewear, facial hair, makeup, non-earlobe piercings, etc.) were not included as experimental targets. To maximize consistency across faces, only photographs of White-appearing individuals who self-identified ages of 18–29 were included. Using Adobe Photoshop CS3 Extended, research assistants then removed the hair and ears from each head and converted the images to grayscale (8-bit bitmap format) to create a standardized 'face' stimulus."

References

- Bailey, J. M., P. L. Vasey, L. M. Diamond, S. M. Breedlove, E. Vilain, and M. Epprecht. 2016. "Sexual Orientation, Controversy, and Science." *Psychological Science in the Public Interest* 17:45–101. <https://doi.org/10.1177/1529100616637616>.
- Bridges, T., and C. J. Pascoe. 2014. "Hybrid Masculinities: New Directions in the Sociology of Men and Masculinities." *Sociology Compass* 8:246–58. <https://doi.org/10.1111/soc4.12134>.
- Bruno, D., M. Lyons, and G. Brewer. 2014. "Response to Bayesian Advice for Gaydar-Based Picking Up: Commentary on Lyons, Lynch, Brewer, and Bruno (2013) by Plöderl." *Archives of Sexual Behavior* 43:11–12. <https://doi.org/10.1007/s10508-013-0222-x>.
- Ceglowski, M. 2016. "The Moral Economy of Tech." Paper presented at the Society for the Advancement of Socio-Economics Annual Meeting, June 26, Berkeley, CA. Retrieved April 25, 2018. http://idlewords.com/talks/sase_panel.htm.
- Cohen, P. 2017. "On Artificially Intelligent Gaydar." *Family Inequality*. <https://familyinequality.wordpress.com/2017/09/11/on-artificially-intelligent-gaydar/>
- Colzato, L. S., L. van Hooijdonk, W. P. M. van den Wildenberg, F. Harinck, and B. Hommel. 2010. "Sexual Orientation Biases Attentional Control: A Possible Gaydar Mechanism." *Frontiers in Psychology* 1:13. <https://doi.org/10.3389/fpsyg.2010.00013>.
- Cox, W. T. L., P. Devine, A. Bischmann, and J. S. Hyde. 2016. "Inferences about Sexual Orientation: The Roles of Stereotypes, Faces, and the Gaydar Myth." *Journal of Sex Research* 53:157–71. <https://doi.org/10.1080/00224499.2015.1015714>.
- Cox, W. T. L., P. Devine, A. Bischmann, and J. S. Hyde. 2017. "Ecological Invalidity of Existing Gaydar Research: In-Lab Accuracy Translates to Real-World Inaccuracy: Response to Rule, Johnson, and Freeman (2016)." *Journal of Sex Research* 54:820–4. <https://doi.org/10.1080/00224499.2017.1278570>.
- DiLallo, K., and J. Krumholtz. 1994. *The Unofficial Gay Manual: Living the Lifestyle (Or at Least Appearing to)*. New York, NY: Doubleday.

- The Economist*. 2017. "Advances in AI Are Used to Spot Signs of Sexuality." *The Economist*, September 9.
- Fasoli, F., and P. Hegarty. 2017. "Straight Talk about Gaydar: How Do Individuals Guess Others' Sexual Orientation?" *The Inquisitive Mind* 7:34.
- Fausto-Sterling, A. 2000. *Sexing the Body: Gender Politics and the Construction of Sexuality*. New York, NY: Basic Books.
- Feltman, R. 2015. "New Study Finds that Your 'Gaydar' Is Terrible." *Washington Post*, September 9.
- France, D. 2007. "The Science of Gaydar." *New York*, June 17.
- Freeman, J. B., K. L. Johnson, N. Ambady, and N. O. Rule. 2010. "Sexual Orientation Perception Involves Gendered Facial Cues." *Personality and Social Psychology Bulletin* 36:1318–31. <https://doi.org/10.1177/0146167210378755>.
- Gelman, A., and D. Park. 2009. "Splitting a Predictor at the Upper Quarter or Third and the Lower Quarter or Third." *American Statistician* 63:1–8. <https://doi.org/10.1198/tast.2009.0001>.
- Hall, M. 2015. *Metrosexual Masculinities*. New York, NY: Palgrave Macmillan. <https://doi.org/10.1057/9781137404749>.
- Hemenway, D. 1997. "The Myth of Millions of Annual Self-Defense Gun Uses: A Case Study of Survey Overestimates of Rare Events." *Chance* 10:6–10. <https://doi.org/10.1080/09332480.1997.10542033>.
- Laumann, E. O., J. H. Gagnon, R. T. Michael, and S. Michaels. 1994. *The Social Organization of Sexuality: Sexual Practices in the United States*. Chicago, IL: University of Chicago Press.
- Lawson, W. 2005. "Queer Eyes: Blips on the Gaydar." *Psychology Today*, November 1.
- Kahneman, D., and A. Tversky. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185:1124–31. <https://doi.org/10.1126/science.185.4157.1124>.
- Magnet, S. A. 2011. *When Biometrics Fail: Gender, Race, and the Technology of Identity*. Durham, NC: Duke University Press. <https://doi.org/10.1215/9780822394822>.
- Mattson, G. 2017. "Artificial Intelligence Discovers Gayface. Sigh." Greggor Mattson Blog. <https://greggormattson.com/2017/09/09/artificial-intelligence-discovers-gayface/>
- Minnelli, L., N. Tennant, C. Lowe, and J. Mendelsohn. 1989. *Results*. New York, NY: Epic Records.
- Moskowitz, P. 2018. "The Faggy Magic of Adam Rippon." *Splinter*, February 16.
- Newport, F. 2015. "Americans Greatly Overestimate Percent Gay, Lesbian in U.S." *Gallup News*, May 21.
- Newton, E. 1984. "The Mythic Mannish Lesbian: Radclyffe Hall and the New Woman." *Signs* 9:557–75. <https://doi.org/10.1086/494087>.
- Olivola, C. Y., and A. Todorov. 2010. "Fooled by First Impressions? Reexamining the Diagnostic Value of Appearance-Based Inferences." *Journal of Experimental Social Psychology* 46:315–24. <https://doi.org/10.1016/j.jesp.2009.12.002>.
- Oxford English Dictionary, Third Edition. 2003. Definition of "Gaydar." Cary, NC: Oxford University Press. Retrieved April 25, 2018 (<http://www.oed.com/viewdictionaryentry/Entry/247143>).
- Plöderl M. 2014. "Bayesian Advice for Gaydar-Based Picking Up: Commentary on Lyons, Lynch, Brewer, and Bruno (2013)." *Archives of Sexual Behavior* 43:7–9. <https://doi.org/10.1007/s10508-013-0178-x>.

- Rieger, G., J. A. Linsenmeier, L. Gygax, S. Garcia, and J. M. Bailey. 2010. "Dissecting 'Gaydar': Accuracy and the Role of Masculinity-Femininity." *Archives of Sexual Behavior* 39:124–40. <https://doi.org/10.1007/s10508-008-9405-2>.
- Saletan, W. 2011. "Read My Lisp." *Slate*, July 15.
- Savin-Williams, R. C. 2006. "Who's Gay? Does It Matter?" *Current Directions in Psychological Science* 15:40–44. <https://doi.org/10.1111/j.0963-7214.2006.00403.x>.
- Schramm, J. D. 2018. "AI 'Gaydar' Could Compromise LGBTQ People's Privacy—and Safety." *Washington Post*, February 19.
- Shelp, S. G. 2003. "Gaydar." *Journal of Homosexuality* 44:1–14. https://doi.org/10.1300/J082v44n01_01.
- Sides, J., and J. Citrin. 2007. "European Opinion about Immigration: The Role of Identities, Interests and Information." *British Journal of Political Science* 37:477–504. <https://doi.org/10.1017/S0007123407000257>.
- Srivastava, S. 2011. "The Precisely Fuzzy Science of Gaydar." *The Hardest Science*. <http://hardsci.wordpress.com/2011/07/21/the-precisely-fuzzy-science-of-gaydar/>
- Tabak, J. A., and V. Zayas. 2012a. "The Roles of Featural and Configural Face Processing in Snap Judgments of Sexual Orientation." *PLoS One* 75:e36671. <https://doi.org/10.1371/journal.pone.0036671>.
- Tabak, J. A., and V. Zayas. 2012b. "The Science of 'Gaydar'." *New York Times*, June 3, SR10.
- Toma, C. L., and J. T. Hancock. 2010. "Looks and Lies: The Role of Physical Attractiveness in Online Dating Self-Presentation and Deception." *Communication Research* 37:335–51. <https://doi.org/10.1177/0093650209356437>.
- Valdes, F. 1995. "Queers, Sissies, Dykes, and Tomboys: Deconstructing the Conflation of 'Sex,' 'Gender,' and 'Sexual Orientation' in Euro-American Law and Society." *California Law Review* 83:12–33. <https://doi.org/10.2307/3480882>.
- Valentova, J. V., K. Kleisner, J. Havlíček, and J. Neustupa. 2014. "Shape Differences between the Faces of Homosexual and Heterosexual Men." *Archives of Sexual Behavior* 43:353–61. <https://doi.org/10.1007/s10508-013-0194-x>.
- Wang, Y., and M. Kosinski. 2018. "Deep Neural Networks Are More Accurate than Humans at Detecting Sexual Orientation from Facial Images." *Journal of Personality and Social Psychology* 114:246–57. <https://doi.org/10.1037/pspa0000098>.
- Zytka, D., S. A. Grandhi, and Q. Jones. 2014. "Impression Management Struggles in Online Dating." Paper presented at the Proceedings of the 18th International Conference on Supporting Group Work, November 9–12, Sanibel Island, FL. <https://doi.org/10.1145/2660398.2660410>.

Acknowledgements: We thank Michal Kosinski and Martin Plöderl for helpful comments.

Andrew Gelman: Department of Statistics and Department of Political Science, Columbia University, New York. E-mail: gelman@stat.columbia.edu.

Greggor Mattson: Department of Sociology and Program in Gender, Sexuality, and Feminist Studies, Oberlin College, Ohio. E-mail: gmattson@oberlin.edu.

Daniel Simpson: Department of Statistical Sciences, University of Toronto, Canada. E-mail: dp.simpson@gmail.com.