

SOCIAL NETWORKS

Leaking privacy and shadow profiles in online social networks

David Garcia

Social interaction and data integration in the digital society can affect the control that individuals have on their privacy. Social networking sites can access data from other services, including user contact lists where nonusers are listed too. Although most research on online privacy has focused on inference of personal information of users, this data integration poses the question of whether it is possible to predict personal information of nonusers. This article tests the shadow profile hypothesis, which postulates that the data given by the users of an online service predict personal information of nonusers. Using data from a disappeared social networking site, we perform a historical audit to evaluate whether personal data of nonusers could have been predicted with the personal data and contact lists shared by the users of the site. We analyze personal information of sexual orientation and relationship status, which follow regular mixing patterns in the social network. Going back in time over the growth of the network, we measure predictor performance as a function of network size and tendency of users to disclose their contact lists. This article presents robust evidence supporting the shadow profile hypothesis and reveals a multiplicative effect of network size and disclosure tendencies that accelerates the performance of predictors. These results call for new privacy paradigms that take into account the fact that individual privacy decisions do not happen in isolation and are mediated by the decisions of others.

INTRODUCTION

The networked nature of our digital society fundamentally changes the principles of how we interact (1). One of these is privacy: Using online services carries privacy losses that are not always trivial to perceive and decide upon, neither for users nor for regulators. Not only does social surveillance allow people to closely watch each other (2), but also the data of an individual user can be used to infer its private attributes (3, 4). From a purely individualist perspective, empowering users to control and price their private information would allow them to balance the benefits, costs, opportunities, and risks of online activity (3, 5). This would hold if individuals used online media in isolation, as it was the case in the early days of the Web. However, the ubiquity of social media renders this individual perspective obsolete and can produce collective effects beyond individual decisions and control (6–10). Users are constantly interacting with each other online, leaving large and deep layers of information that can reveal private attributes of others without their awareness (11). It is possible that the control of individuals over their information is progressively being lost through leaking privacy, leaving a trace of private information with each social interaction.

An example of leaking privacy is the phenomenon of shadow profiles: files with private information of a person that online services can generate from the data that the social contacts of that person give to the service (12). Shadow profiles could be constructed without permission or knowledge of the person who is being profiled, who might not be a user nor agree to the terms of the online service that builds the profile. The idea of shadow profiles came to light in 2013, when a bug in Facebook revealed that the mobile phone numbers of some users had been extracted from the phonebooks of their friends but never provided by the users themselves (13). Because many online services have access to user contact lists outside the service, for example, Facebook's messenger phone app permissions and its potential connection with WhatsApp, the same inference of personal information could be carried

out for people who are not users. To ensure the right to privacy and informational self-determination (14), we need to evaluate whether shadow profiles are a possibility. This question is formalized as the shadow profile hypothesis: The data given by users of an online service predict personal information of nonusers.

Previous research on privacy in social media provides background on the inference of private attributes of users from their online activity (15). Some examples of this line of research are the prediction of gender, age, and political orientation with Twitter data (16), and of sexual orientation and romantic partnerships with Facebook data (17, 18). These predictions build on the information captured by assortativity and homophily in social networks (16), providing evidence that private attributes of users can be predicted when sufficient contextual data are available. These analyses evaluate how some information about a user can be predicted through its activity and the activity of its friends but do not venture to evaluate whether these predictions can be applied to people who are not users of the service. Notable exceptions have applied simulation approaches to investigate the inference of friendships outside Facebook (19) and used friendship signals to infer sexual orientation (9), but we lack an empirical and formal test of the shadow profile hypothesis in a large online social network.

The research gap that has so far prevented the analysis of predictive power over nonusers can be explained by the lack of necessary data outside the private control of the owners of online services. A large company, such as Facebook or Google, could publicly show the possibility to build shadow profiles, but these results could easily be in conflict with the company's interests and business models. Therefore, we need audits by independent researchers to reliably test whether shadow profiles are a possibility. To overcome this challenge and provide a first test of the shadow profile hypothesis, we use the method of Internet Archaeology (8): We study the traces of a disappeared online social network, Friendster, to address a question about its functioning. Here, we test the shadow profile hypothesis against the data that were abandoned in Friendster when it was discontinued as an online social network but captured by the Internet Archive and made available for independent research. We trace back the history of the growth of the

Copyright © 2017
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
NonCommercial
License 4.0 (CC BY-NC).

social network to evaluate whether information inside the network had predictive power to infer personal information of individuals who were not users at that time, with the aim to empirically test the shadow profile hypothesis.

We apply principles from network science to gain insight into the structural properties that can explain whether shadow profiles can be built, measuring how personal attributes of users are related to their neighbors. We evaluate a straightforward prediction method for private information on the basis of the data of the friends of a user to then historically evaluate that prediction for nonusers as the network grew. The aim of this article was not to advance the techniques to infer personal information of individuals outside a social network but to measure a lower bound on that predictability and use it to test the shadow profile hypothesis.

RESULTS

Neighborhood mixing patterns of user profiles

We can observe the mixing patterns of personal information in the neighborhood vectors of users with different sexual orientations and relationship status. Figure 2 shows radar plots with the normalized log frequency of neighbors of each class as a function of the class of the individual user (see Materials and Methods for details about the calculation). Sexual orientation displays mixing patterns: Heterosexual users are less likely to be connected to users of the same gender; bisexual users are more likely to be connected to other bisexual users of any gender; and homosexual users are strikingly more likely to be connected to homosexual users of the same gender.

The mixing patterns of sexual orientation are more complex than assortativity: Whereas homosexual users tend to be connected to other users of the same orientation and gender, the pattern for heterosexual users is of heterophily with respect to gender, that is, they are more likely to connect to users of the opposite gender. On the contrary, relationship status displays a pattern more typical of assortativity. All five classes of relationship status have higher probability to be connected to users with the same status, and this pattern is the strongest for users with a status of married and domestic partners.

It is worth noting that the mixing patterns of sexual orientation can be noticed at longer neighborhood distances. Text A (Supplementary Materials) presents the radar plots and assortativity mixing matrices for neighborhoods at distances 2 and 3. The same patterns at distance 1 can be observed at distance 2 with weaker strength, and only the assortative mixing of homosexual users can be appreciated at distance 3. This points that user information beyond the local neighborhood of a user could be predictive of personal information, but in this article, we focus on metrics at distance 1 as a lower bound to predictor performance.

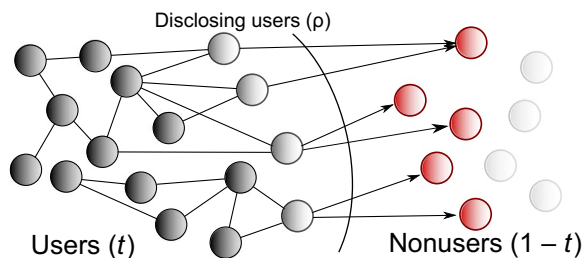


Fig. 1. Shadow profile problem. Diagram of the shadow profile problem, where the network contains a fraction t of the final size and users in the network have a tendency r to share their contact lists.

Social inference of profile attributes

We define the prediction problem as a binary classification as explained in Materials and Methods, evaluating the performance of unsupervised neighbor frequency-based predictors. We first explore the predictive power that social network data have for personal information of users inside the network to test the same case for nonusers in the next section. The first two panels in Fig. 3 show the predictor performance for relationship status and sexual orientation using all information available in the network. The receiver operating characteristic (ROC) curves are significantly above the diagonal dashed lines, with AUC (area under the curve) values above 0.6. We repeat the same prediction task over a 100-fold evaluation in subsets of 1% of the data and got similar values significantly above 0.5. Furthermore, a repetition of the prediction with permuted sexual orientation classes shows an AUC extremely close to 0.5, clarifying that the above result is not a false positive.

Although the above analysis shows evidence of the informativeness of neighborhood information, not all users share their private information in the social network. We analyze how the general tendency to disclose private information affects the prediction performance. We quantify the disclosure tendency δ as the probability of each user sharing its personal information and repeat the prediction task when sampling users according to that tendency (see Materials and Methods). The left panel in Fig. 3 shows the average predictor performance for both private attributes in 100 samples given each disclosure tendency value between 0.1 and 0.9. The monotonically increasing pattern is evident: The higher the tendency of users to disclose personal information in the social network, the better the prediction that could be carried out for those users who chose not to share that information.

Testing the shadow profile hypothesis

We test the shadow profile hypothesis in an audit that evaluates whether the prediction performance of personal data of nonusers improved as the network grew and users shared their contact lists. The problem, explained in detail in Materials and Methods, is outlined in Fig. 1. At a point in time when a fraction t of all the users were inside (and a fraction $1 - t$ were nonusers), each user inside the network shared their contact lists with a probability ρ . The users sharing these contact lists are disclosing users, and through their contact lists, it is possible to make a prediction of personal information of nonusers. We measure the performance of the predictors of sexual orientation and relationship status in 100 samples of disclosing users for each value of ρ between 0.1 and 1 by increments of 0.1 and over the growth of the social network from $t = 0.1$ to $t = 0.9$ by increments of 0.1.

The performance of the predictor monotonically increases with t and ρ . Figure 4 shows the mean performance over 100 samples for the prediction of sexual orientation and relationship status of nonusers. The increasing pattern of AUC with t and ρ is evident, approaching values above 0.57 for sexual orientation and 0.62 for relationship status. To formally test the shadow profile hypothesis, we compute the Kendall τ correlations between AUC, t , and ρ . As shown in Table 1, predictor performance is positively correlated with t and ρ for both sexual orientation and relationship status, lending strong evidence that supports the shadow profile hypothesis.

To understand the combination of roles of t and ρ in predictor performance, we computed partial Kendall correlation coefficients (explained in Materials and Methods). Partial correlation coefficients, as shown in Table 1, are strong and significant in both cases, indicating that both parameters have a positive effect on AUC. This effect is

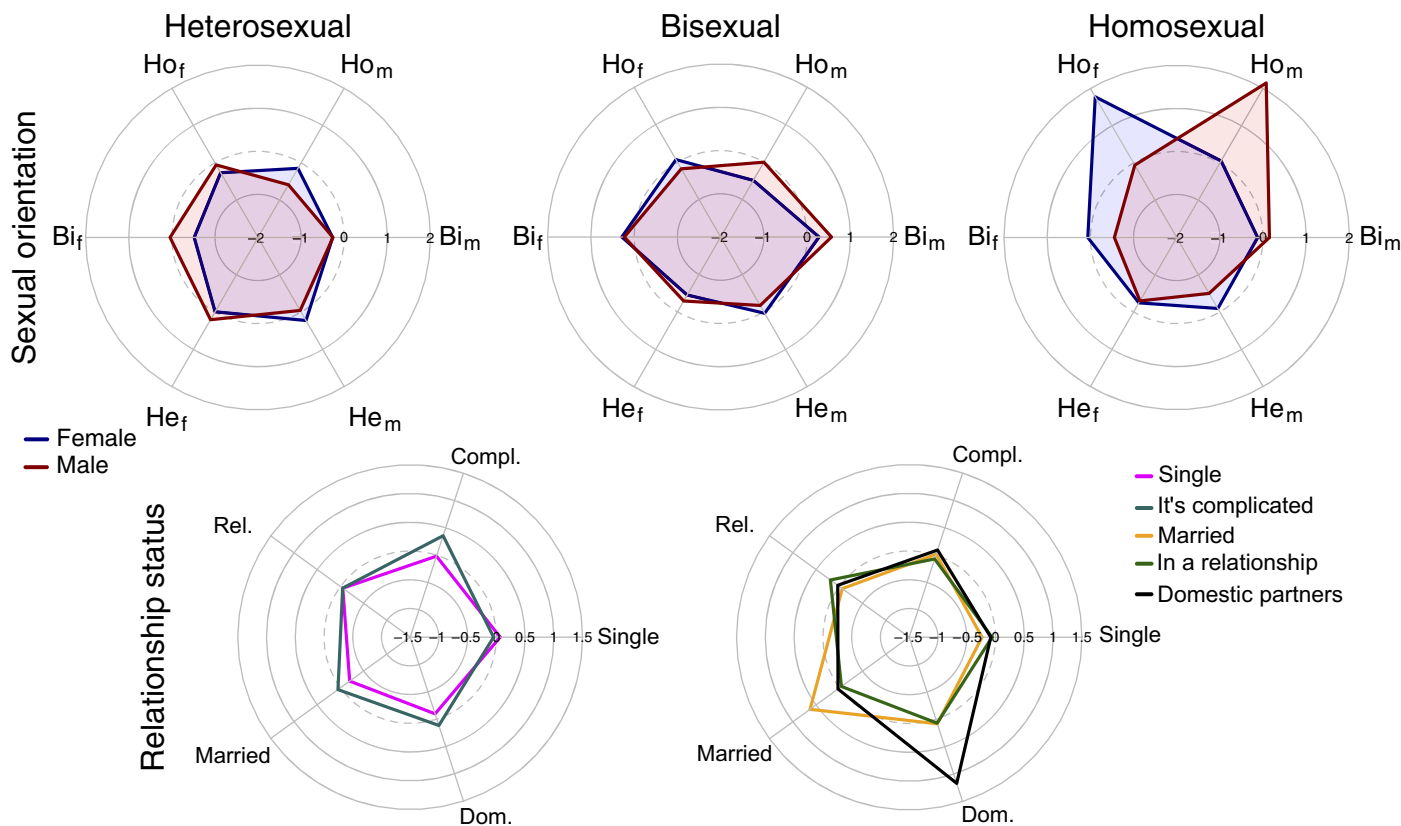


Fig. 2. Radar plots of sexual orientation and relationship status in individual neighborhoods. Each point in a radar plot has a distance from the center corresponding to the logarithm of the normalized ratios of neighbors of each orientation and relationship status. Each radar plot corresponds to one class of relationship status or sexual orientation. Relationship status displays a pattern of assortativity with the same status, and sexual orientation displays a more complex mixing pattern with homophily for homosexual users and heterophily for gender for heterosexual users. Ho_f , homosexual female; Ho_m , homosexual male; Bi_f , bisexual female; Bi_m , bisexual male; He_f , heterosexual female; He_m , heterosexual male.

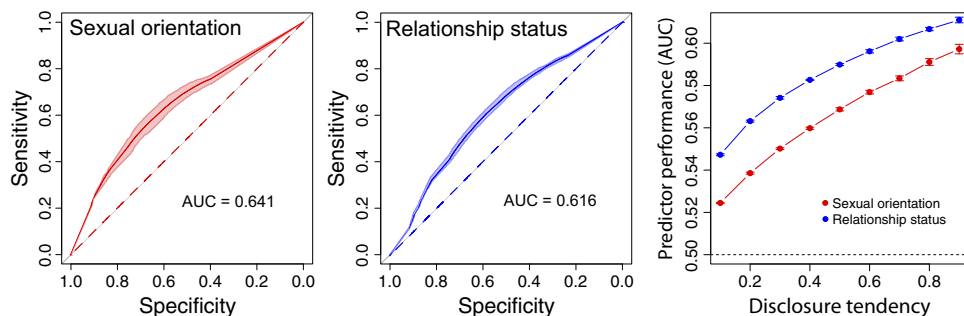


Fig. 3. Predictor performance inside the social network. ROC curves of the predictors of sexual orientation (left) and relationship status (middle). The shaded areas show the maximum and minimum values of the predictor performance in a 100-fold evaluation. The diagonal dashed line is the result of the same analysis with permuted profiles [area under the curve (AUC) = 0.499]. In both cases, neighborhood information in the social network predicts individual information. Right: AUC for various values of the disclosure tendency to share personal information publicly in the social network. Error bars show SD of 100 samples given the value of the disclosure tendency. AUC increases with disclosure tendency in both cases.

multiplicative; the correlation coefficient of AUC with the product of t and ρ is even stronger, reaching values above 0.9. This observation is robust when computing it as a partial correlation with t and ρ . This indicates that, if users have a higher tendency to disclose their contact lists, this has a multiplier effect over network growth toward high predictor accuracy. All the results of Kendall correlation coefficients are significant in bootstrapping and permutation tests, as reported in text B

(Supplementary Materials), revealing the robustness of the evidence that supports the shadow profile hypothesis.

DISCUSSION

The results of this analysis indicate that personal information of people outside Friendster could have been predicted with the data

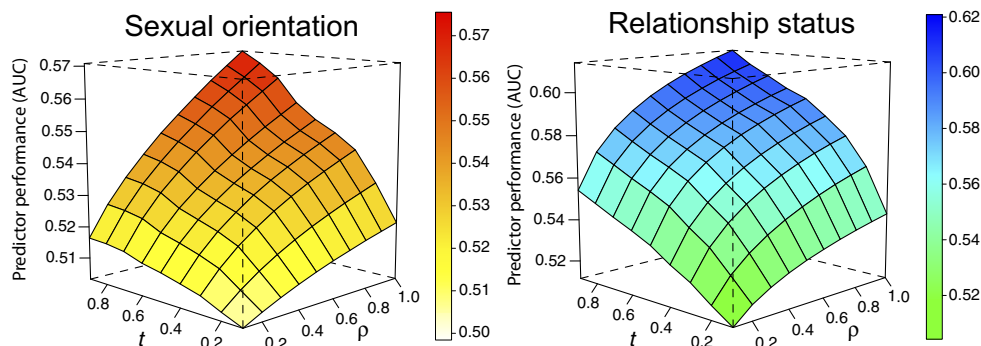


Fig. 4. Shadow profile predictor performance. Predictor performance (AUC) versus t and ρ for sexual orientation and relationship status. Each value is computed as the mean of 100 samples for each combination of t and ρ values. In both cases, predictor performance increases with network size and tendency to share contact lists.

Table 1. Kendall correlation coefficients of predictor performance (AUC) versus t and ρ . Estimates are median values of bootstrap distributions and confidence intervals are calculated at the 95% level. P values result from permutation tests with 10,000 permutations of the AUC values. AUCs of both sexual orientation and relationship status increase monotonically with t and ρ . This effect is multiplicative, reaching high τ values for the correlation between AUC and the product of t and ρ . All these observations are robust to the same tests with partial correlation coefficients and significant with $P < 0.05$ in permutation tests.

Coefficient	Sexual orientation			Relationship status		
	Estimate	Confidence interval	P	Estimate	Confidence interval	P
$\tau(\text{AUC}, t)$	0.476	0.465–0.488	<0.05	0.582	0.573–0.592	<0.05
$\tau(\text{AUC}, \rho)$	0.569	0.558–0.578	<0.05	0.464	0.453–0.475	<0.05
$\tau(\text{AUC}, t\rho)$	0.951	0.949–0.952	<0.05	0.941	0.940–0.943	<0.05
$\tau(\text{AUC}, t \rho)$	0.579	0.573–0.585	<0.05	0.657	0.652–0.663	<0.05
$\tau(\text{AUC}, \rho t)$	0.647	0.641–0.653	<0.05	0.571	0.565–0.577	<0.05
$\tau(\text{AUC}, t\rho t)$	0.936	0.935–0.938	<0.05	0.918	0.917–0.920	<0.05
$\tau(\text{AUC}, t\rho \rho)$	0.930	0.928–0.932	<0.05	0.926	0.924–0.927	<0.05

shared by Friendster users, supporting the shadow profile hypothesis. This phenomenon can be explained by the mixing patterns of sexual orientation and relationship status that existed in the network, which revealed information of nonusers through the contact lists of users. The historical analysis of how predictor performance depends on the size of the network and disclosure tendencies sheds light on how these two factors interact. They do not simply add up in predictive power; hence, they display a multiplicative effect by which the marginal effect of one factor increases with the other. Permutation, bootstrapping, and partial correlation tests show the robustness of these results, evidencing that social networking sites are not powerless in the task of creating shadow profiles of nonusers.

All classifiers had moderate predictor performances, with AUC values around 0.6. The relatively poor performance of these predictors is a result of the straightforward and unsupervised method applied here, which is based on a one-step rule over the neighborhood of users. The aim of this article was to evaluate whether predictor performance for nonusers grows with the amount of data in the network. To prevent any ethical issue regarding private information, the complexity of the predictor was kept low on purpose, and hence the moderate predictor performance. This strategy enabled the measurement of a lower bound on predictor performance without advancing the methods to construct shadow profiles, which could pose a threat to privacy if ap-

plied to other social media and private attributes. This approach has proved accurate enough to be significantly better than random and sensitive enough to test the shadow profile hypothesis.

The most important limitation of this article is the external validity of historical and observational data analysis. The results presented here used the data of future users as ground truth, and thus, it is open for future research to test how the shadow profile hypothesis can be tested with data exogenous to online social networks. We do not have reason to assume that Friendster had better data than current social networking sites, but it is nevertheless necessary to replicate the results of this article with online services that are currently in operation. Survey approaches can gather data from volunteering users (3, 4) and nonusers and, thus, provide an in vivo test of the shadow profile hypothesis that can evaluate to which extent private information of nonusers can be predicted today. Furthermore, independent data audits on the information workflows of online services can replicate the principles of this article and measure in real time whether personal information of nonusers is being directly or indirectly inferred.

Keeping information private was possible by physical means decades ago, hence the term private space. However, the widespread adoption of information and communication technologies turns our information into a much more pervasive and diffusive element that cannot be conceived as a space or material that we can directly control.

The persistent traces of our online social interaction can slowly accumulate enough data to effectively diminish the decision power of an individual to keep personal information private. In this leaking privacy scenario, our private information resembles a liquid that is slowly shared with others in our online activity and can be analyzed without the control or permission of the owner of the information. Regulations to protect the right of individuals to decide about their personal information need to take into account this new phenomenon, avoiding individualized paradigms that ignore the power of social inference and the evident possibility for social networking sites to construct shadow profiles.

MATERIALS AND METHODS

Experiment design

When Friendster announced the discontinuation of their social networking services in 2011, the Internet Archive retrieved all publicly available information on more than 100 million accounts (20). Previous research on this data set shows that the first 20 million user accounts were created in wave of early user adoption in the United States, who later left for other social networking sites such as Facebook and Myspace (8). We used this set of early accounts for our analysis, quantifying personal information as the social network grew from its creation. The Internet Archive data set contains Web snapshots of user profiles that contain friend lists and personal information such as gender, relationship status, and romantic interests. Of the first 20 million accounts, 6 million had public friendship lists, and 3.3 million had public personal information. From all the information available on the profiles, we processed the personal attributes of relationship status, gender, and romantic interests. Many users reveal the genders they are romantically interested in, from which we can classify the possible sexual orientations in the data set: bisexual female, bisexual male, heterosexual female, heterosexual male, homosexual female, and homosexual male.

Ethical considerations

This research uses only publicly available information and has not access to any other source besides the Internet Archive data set. The analysis system anonymized any identifiable information and used only obfuscated identifiers (IDs) to analyze the social network. This research constitutes an archival study that (i) does not interact with any individual user, (ii) only uses information consciously made public by users, and (iii) does not infer any new personal information; thus, this study is exempt of a review from an institutional review board. The exercise to test the shadow profile hypothesis is a post hoc audit of how the owners of the social networking site could have predicted personal information that was private to them at a point in time but does not improve prediction methods nor infer information that is private to us at the time of analysis. This research provides a data-driven audit to make better informed decisions about privacy in the future, a benefit that greatly outweighs the minimal risks of analyzing archival online user data.

The shadow profile problem

We formalized the problem of predicting sexual orientation and relationship status by binarizing both attributes before analysis. For the case of sexual orientation, we classified as $O_i = 1$ every user i that is either homosexual or bisexual regardless of gender (that is, a lesbian, gay, or bisexual user) and as $O_i = 0$ if the user is heterosexual of any gender. For relationship status, we classified as $R_i = 1$ if user i declares

to be in some sort of relationship (married, domestic partners, or in a relationship) and $R_i = 0$ if the user has the status of single or it's complicated. Given this binarization, each user i has two neighborhood ratios: $o_i = \frac{1}{N_i} \sum_{j \in \Gamma_i} O_j$ and $r_i = \frac{1}{N_i} \sum_{j \in \Gamma_i} R_j$, where Γ_i is the set of friends of user i and N_i is the amount of friends of i . From these ratios, we defined the straightforward, unsupervised one-step predictors $\hat{O}_i = \Theta[O_i - T_O]$ and $\hat{R}_i = \Theta[R_i - T_R]$, where $\Theta[x]$ is the Heaviside step function that takes the value of 1 if $x > 0$ and 0 otherwise, and T_O and T_R are the sensitivity thresholds used to compute ROC curves as explained below.

We evaluated the performance of the above predictor in three scenarios. First, we used all available network data to produce an estimate of the upper bound of predictor performance on a 100-fold evaluation over a random partition of the data in subsets of 1% of the total amount of users. Second, we defined a disclosure tendency parameter δ , which measures the probability of an individual user to share its personal information in the social network. Given a value of δ , we generated 100 samples of the network in which the personal data of each user are included according to an independent Bernoulli trial with probability δ . Over each of these samples, we performed the above prediction with all available data to measure the relationship between predictor performance and disclosure tendency.

In the third prediction scenario, we historically evaluated the problem as shadow profile generation. We defined a point in time in which a fraction t of the users had joined the network, and the rest did not join yet (we can identify this given the user ID sequence of the data set). At this point in time, the ground truth for our tests is the personal data of future users, given that all of them were not sharing personal data with Friendster. In addition, we defined the tendency of users in the network to share their full contact lists as ρ . We created 100 samples in which each user has shared their contact list according to Bernoulli trials with probability ρ and denoted those users as disclosing users. In this way, given a value of t , a value of ρ , and a sample of disclosing users sharing their contact lists, we can perform and evaluate a prediction, as described in Fig. 3. This prediction uses the same rules as above, calculating the ratios o_i and r_i of nonusers over the set of their friends that are disclosing users inside the network.

Statistical analysis

To understand the mixing patterns of the whole network, we calculated neighborhood vectors for each user. Given a user i , its neighborhood vector of sexual orientation contains an entry with the logarithm of the ratio between the fraction of neighbors of each sexual orientation in the neighborhood of i and the total fraction of users with that sexual orientation in the whole network. In this way, each entry has a positive value if a certain class of users is overrepresented in the neighborhood of user i and a negative value if they are underrepresented. We similarly constructed the neighborhood vector of relationship status, having this way a total of 11 entries in two vectors for each user. We used these values to analyze neighborhood mixing patterns.

We evaluated the predictors \hat{O}_i and \hat{R}_i by computing the ROC curves (21) with the pROC R package (22), calculating sensitivity and specificity over all possible values of the thresholds T_O and T_R . We quantified predictor performance as the AUC of the ROC, which takes a value of 0.5 for uninformative predictors and 1 for perfect predictors (23).

For the evaluation of the predictor performance in the generation of shadow profiles given a value of t and ρ , we computed the AUC of \hat{O}_i and \hat{R}_i over 100 samples of disclosing users. We iterate t from 0.1 to

0.9 by increments of 0.1 and ρ from 0.1 to 1 by increments of 0.1. We tested the shadow profile hypothesis, that is, that AUC grows with t and ρ , by computing the Kendall correlation coefficients $\tau(\text{AUC}, t)$ and $\tau(\text{AUC}, \rho)$ (24). To test that both factors have independent effects, we computed partial Kendall correlation coefficients with the ppcor R package (25): $\tau(\text{AUC}, t | \rho)$ and $\tau(\text{AUC}, \rho | t)$. Furthermore, to evaluate whether t and ρ display a multiplicative effect, we calculated the correlation between predictor performance and their product $\tau(\text{AUC}, t \cdot \rho)$ and tested its robustness with the corresponding partial correlation coefficients $\tau(\text{AUC}, t \cdot \rho | \rho)$ and $\tau(\text{AUC}, t \cdot \rho | t)$. We tested the statistical significance of all these correlation coefficients in two ways. First, we computed 10,000 bootstrap estimates of each correlation coefficient. Second, we performed a permutation test by computing the same correlation coefficients over 10,000 permutations of predictor performance for each correlation coefficient. The bootstrap estimates allow us to measure the 95% confidence interval of each coefficient, and the permutation test allows us to measure a P value against the null hypothesis of an absence of correlation. The results of all these tests are reported on text B and allow us to have a robust nonparametric test of the shadow profile hypothesis.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/3/8/e1701172/DC1>

Text A: Mixing patterns of sexual orientation at distances 2 and 3

Text B: Prediction performance details

fig. S1. Radar plots of sexual orientation at neighborhoods at distances 2 and 3.

fig. S2. Mixing patterns of sexual orientation at distances 1 to 3.

fig. S3. Predictor performance versus t and ρ .

fig. S4. Bootstrapping and permutation tests of the correlation between sexual orientation AUC and t , ρ , and their product.

fig. S5. Bootstrapping and permutation tests of the correlation between relationship status AUC and t , ρ , and their product.

fig. S6. Bootstrapping and permutation tests of the partial correlation between sexual orientation AUC and t , ρ , and their product.

fig. S7. Bootstrapping and permutation tests of the partial correlation between relationship status AUC and t , ρ , and their product.

REFERENCES AND NOTES

1. M. Castells, *The Rise of the Network Society* (Blackwell Publishers, 2000).
2. A. Marwick, The public domain: Surveillance in everyday life. *Surveillance & Society* **9**, 378–393 (2012).
3. M. Kosinski, D. Stillwell, T. Graepel, Private traits and attributes are predictable from digital records of human behavior. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 5802–5805 (2013).
4. W. Youyou, M. Kosinski, D. Stillwell, Computer-based personality judgments are more accurate than those made by humans. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 1036–1040 (2015).
5. Y. Liu, K. P. Gummadu, B. Krishnamurthy, A. Mislove, Analyzing Facebook privacy settings: User expectations vs. reality, in *Proceedings of the 2011 ACM SIGCOMM Conference on Internet Measurement Conference*, 2 to 4 November 2011, pp. 61–70.
6. J.-P. Onnela, F. Reed-Tsochas, Spontaneous emergence of social influence in online systems. *Proc. Natl. Acad. Sci. U.S.A.* **107**, 18375–18380 (2010).
7. W. Lance Bennett, A. Segerberg, The logic of connective action: Digital media and the personalization of contentious politics. *Information, Communication & Society* **15**, 739–768 (2012).
8. D. Garcia, P. Mavrodiev, F. Schweitzer, Social resilience in online communities: The autopsy of Friendster, in *Proceedings of the First ACM Conference on Online Social Networks*, 7 to 8 October 2013, pp. 39–50.
9. E. Sarigol, D. Garcia, F. Schweitzer, Online privacy as a collective phenomenon, in *Proceedings of the Second ACM Conference on Online Social Networks*, October 1 to 2 2014, pp. 95–106.
10. J. Kleinberg, Analysis of large-scale social and information networks. *Philos. Trans. A Math. Phys. Eng. Sci.* **371**, 20120378 (2013).
11. D. Boyd, Networked privacy. *Surveillance Soc.* **10**, 348–350 (2012).
12. K. Knibbs, What's a Facebook shadow profile and why should you care? *Digital Trends* <http://bit.ly/2oaoTFE> (2013).
13. V. Blue, Anger mounts after Facebook's 'shadow profiles' leak in bug. *ZDNet* <http://zd.net/2o5YhL2> (2013).
14. A. Rouvroy, Y. Poulet, The right to informational self-determination and the value of self-development: Reassessing the importance of privacy for democracy, in *Reinventing Data Protection?* S. Gutwirth, Y. Poulet, P. De Hert, C. de Terwangne, S. Nouwt, Eds. (Springer, 2009), pp. 45–76.
15. E. Zheleva, L. Getoor, To join or not to join: The illusion of privacy in social networks with mixed public and private user profiles, in *Proceedings of the 18th International Conference on World Wide Web*, April 20 to 24 2009, pp. 531–540.
16. F. Al Zamil, W. Liu, D. Ruths, Homophily and latent attribute inference: Inferring latent attributes of Twitter users from neighbors, in *Proceedings of the Sixth International Conference on Weblogs and Social Media*, 2012, pp. 387–390.
17. C. Jernigan, B. F. T. Mistree, Gaydar: Facebook friendships expose sexual orientation. *First Monday* **14** (2009).
18. L. Backstrom, J. M. Kleinberg, Romantic partnerships and the dispersion of social ties: A network analysis of relationship status on Facebook, in *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work and Social Computing*, 15 to 19 February 2014, pp. 831–841.
19. E.-Á. Horvát, M. Hanselmann, F. A. Hamprecht, K. A. Zweig, One plus one makes three (for social networks). *PLOS ONE* **7**, e34740 (2012).
20. Internet Archive, Friendster social network dataset: Friends (2011); <https://archive.org/details/friendster-dataset-201107>.
21. T. Fawcett, An introduction to ROC analysis. *Pattern Recogn. Lett.* **27**, 861–874 (2006).
22. X. Robin, N. Turck, A. Hainard, N. Tiberti, F. Lisacek, J.-C. Sanchez, M. Müller, pROC: An open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinformatics* **12**, 77 (2011).
23. J. A. Hanley, B. J. McNeil, A method of comparing the areas under receiver operating characteristic curves derived from the same cases. *Radiology* **148**, 839–843 (1983).
24. M. G. Kendall, *Rank Correlation Methods* (Charles Griffin & Co., 1948).
25. S. Kim, ppcor: An R package for a fast calculation to semi-partial correlation coefficients. *Commun. Stat. Appl. Methods* **22**, 665–674 (2015).

Acknowledgments: D.G. would like to thank E. Sarigol for technical support and F. Schweitzer for useful discussions. D.G. thanks the Internet Archive for collecting and distributing the data set. **Funding:** This research has been funded by the ETH Foundation through the ETH Risk Center Seed Project "Systemic Risks for Privacy in Online Interaction." **Author contributions:** D.G. designed research, performed analysis, and wrote the article. **Competing interests:** The author declares that he has no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials. All data analyzed in this article are available on the Internet Archive: <https://archive.org/details/archive-team-friendster>.

Submitted 17 April 2017

Accepted 29 June 2017

Published 4 August 2017

10.1126/sciadv.1701172

Citation: D. Garcia, Leaking privacy and shadow profiles in online social networks. *Sci. Adv.* **3**, e1701172 (2017).