Christian Czarnecki et al. (Hrsg.): Workshops der INFORMATIK 2018,
Lecture Notes in Informatics (LNI), Gesellschaft für Informatik, Bonn 2018 109

I know who you are: Deanonymization using Facebook Likes

Sylvio Rüdian[1], Niels Pinkwart[2] und Zhi Liu[3]

Abstract: This paper presents a method to deanonymize people using fanpages' Likes of Facebook users. The strategy shows that information of Likes can be easily crawled from Facebook. Combined with an interactive version of browser-history-stealing it can be used to get identities of users on a website. The attack is possible because of the existence of Facebook' Likes that can be used as a fingerprint. The claim was tested and discussed with real-world collected data. With the assumption of at least 4 collected Likes per user, 99.91% of them can be deanonymized through the fingerprint of Likes. Apart from that we provide potential solutions for protection of identities in social media.

Keywords: social media, deanonymization, browser-history-stealing, fake-captcha, Facebook, Likes, fingerprint, privacy, attack

1 Introduction

Social media like Facebook is one of the most widely used channel of people's worldwide communications. People share much information about their personality at their social media profiles. Simultaneously, there is a missing understanding about issues like privacy. A study has shown that "privacy settings match users' expectations [in] only 37% of the time" [LI11]. Although Facebook has implemented mechanisms for data protection, an attacker still can get much personal data that is accessible through the social network. If these data will be combined with data of a third party, it is possible to identify users by using their public information.

The paper is organized as follows. In section II we summarize related work regarding deanonymization attacks. Section III gives a formal technical definition of the general deanonymization approach and describes the technical methodology how to collect Likes data from Facebook users. Section IV shows the results of our experiments and V proposes some solutions for protection from different points of view. Due to the in-depth privacy protection in Germany, this paper focuses on findings related to Germany.

[1] Humboldt-Universität zu Berlin, Department of Computer Science, Berlin, Germany, ruediasy@informatik.hu-berlin.de

[2] Humboldt-Universität zu Berlin, Department of Computer Science, Berlin, Germany, niels.pinkwart@hu-berlin.de

[3] Central China Normal University, National Engineering Research Center for E-Learning, Wuhan, China, zhiliu@mail.ccnu.edu.cn

2 Related Work

In 2011, Weinberg et al. [WE11] developed an attack to show that memberships of users belonging to groups can be used as a fingerprint. By a prepared website, the researchers were able to get the information of previous visited websites using techniques of browser-history-stealing [RU14]. They assumed that members of groups also visited these groups with their browsers before so that the links of groups can be explored in the users' browser histories. Additionally, an attacker could use browser-history-stealing to determine which groups a user has visited before. With that approach, the researchers were able to deanonymize 22.9% of Facebook users.

Narayanan et al. used the structure of the social media graph to deanonymize people. Researchers used the knowledge of the overlap between different social networks and created an algorithm to "successfully de-anonymize several thousand users in the anonymous graph" of social networks [SH08]. The researchers also have investigated into deanonymization of a Netflix dataset [NA08]. They had shown that it was possible to derive identities from anonymized datasets. With "contextual and background knowledge, as well as cross-correlation with publicly available databases [they were able] to re-identify individual data records" [NA08]. Ji et al. [JI14] investigated into deanonymization of structural data that had been published. They created a general data model for deanonymization and also conducted experiments with social networks.

3 Methodology

3.1 Theoretical approach

The following section describes the approach of deanonymization in a theoretical way by using mathematical sets. In general, deanonymization is the process where firstly data points of several people were collected. After the process of data collection, an attacker tries to find specific values of data points by using a third party, e.g. a website. Values of data points can be combined with the already known dataset of the attacker. Then he has the possibility to derive the identity of persons, although he did not get this information from the third party. To formalize this approach, the dataset C consists of different people P_i with $i \in \mathbb{N}\ \ , i \leq n, n = |P|$ and a set D, where d_j with $j \in \mathbb{N}\ \ , j \leq m,\ m = |D|$ are specific data points in D, e.g. IDs of groups. There are some connections between P and D, defined as $R{:}\, P \times D$ with $(P_i, d_j) \in R$. This can be memberships of persons to groups. Each person P_i can be connected with different data points d_j and each data point d_j can be connected with different persons P_i. For each person P_i there is a set of $C \subseteq D$, where connections to d_j are existing. This tuple C is called fingerprint of person P_i. The attacker can use a third party like a website where he wants to identify persons within the dataset, although he does not know the identity yet. Therefore, he has to find out whether the different data points of D belong to P_i, e.g. whether the current person is a member of several groups. It can be tested with hacking techniques or indirect requests to the user. This results in a new dataset A, where A only contains data points d_j that belong to the concrete person P_i and where at least one connection R exists. Ideally, $A = C$ where C is the fingerprint of person P_i. While the number of hits will be increased, A becomes more detailed to derive the identity using R. Fig. 1 shows an example where P_2 could be deanonymized by the knowledge of two hits d_1 and d_2 (e. g. being member of groups d_1 and d_2) using a prepared website, where $\{d_1, d_2\} = C$ and the user is a member of groups d_1 and d_2.

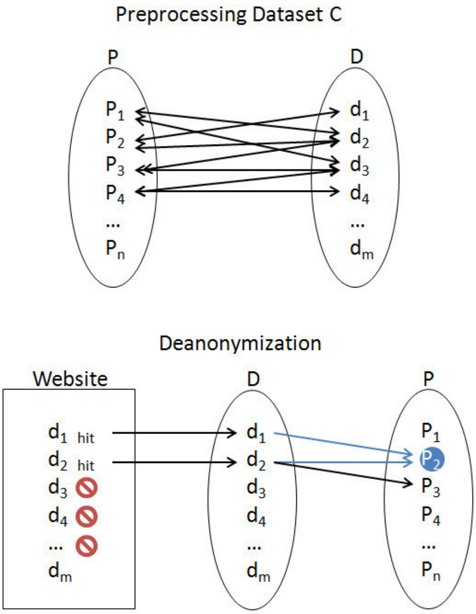

Fig. 1: Preprocessing of the attackers dataset and deanonymization attack to derive an identity.

Weinberg et al. [WE11] focused on memberships to groups as a fingerprint because memberships of public groups are easy to access, without any limitations. Fig. 2 shows the different amount of members of groups (blue) and fans of fanpages (red) for a sample of 100,000 fanpages and 100,000 public groups. It shows that most fanpages have a major amount of connections to users than to groups. “Likes represent one of the most generic kinds of digital footprint. For instance, liking a brand or a product offers a proxy for consumer preferences and purchasing behavior; music-related Likes reveal music taste; and liked web-sites allow for approximating web browsing behavior.” [YO15]

In Germany, users like 28 fanpages on average [SO13]. Worldwide, the average user likes 40 pages [SO13]. At the same time, each user is only a member of 12 groups on average [SA11]. This fact leads to the assumption that Likes are better qualified than memberships to groups since more data (Likes) lead to a fingerprint that can be much more detailed. An attacker can collect a huge amount of Facebook' Likes of fanpages. We assume that one user which likes fanpages also visited the fanpages with his browser before.

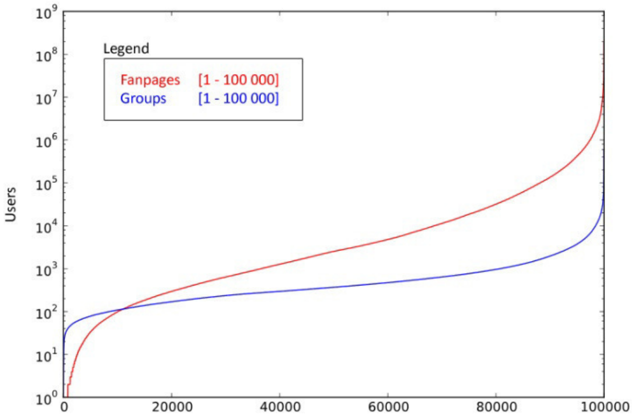

Figure 2: Comparison of the members of 100.000 Facebook' groups and 100.000 fanpages, ordered by amount.

In this paper we use Facebook users' Likes for deanonymization. We aim to answer the question whether "Likes" of fanpages can be used as a fingerprint and how this information can be used to identify peoples' Facebook profiles in a real world scenario. To do that, we prepared a website and used the browser-history-stealing attack to determine previous visited fanpages. The information of previous visited fanpages can be combined with the attackers' dataset to derive identities. The following section describes our approach of collecting data in detail. Additionally, we discuss how people can protect themselves in general and which possibilities browsers and social networks have to prevent their users from the deanonymization.

3.2 Crawling users' profiles and Fanpages

First, we crawled Facebook and stored the connections between users and Likes of Fanpages. Next to the regular website, Facebook also offers a version of the social network for older browsers or Smartphone's. It can be accessed by the URL
`https://mbasic.facebook.com`
This website can be easily used by a crawler [SE17]. After collecting public groups, an attacker can download a list of each groups' members. This can be done with a regular Facebook login, which the crawler uses to get access to the social network. Since Facebook uses predictable URLs for groups, a crawler could download the list of members by downloading the URL

`https://m.facebook.com/groups/[group-ID]` for each known group, where the group-IDs from the last step could be used. The downloaded lists contain usernames of members. If usernames are known, public Likes can be collected by crawling the following URL: `https://mbasic.facebook.com/[username] ?v=likes`

This leads to a list of previously liked fanpages of the user and can be stored. It is only accessible if the current profile is public or if there is an existing friendship relationship between the crawler and the considered profile. A single crawler can get the Likes of nearly 10,000 users after the crawler will get a restricted access for one week (7 days).

Another approach to get the Likes of users can be achieved by using fanpages' lists of users. According to Facebook, these lists only exist for administrators owned fanpages: "You cannot get a list of all the fans of a Facebook Page" [FA14]. The fanpage-id can be extracted by the Facebook-API. A crawler could visit the following URL, to get the list of users that liked the fanpage with a given fanpage-id, although he is not an administrator of the fanpage: `https://www.facebook.com/search/[Fanpage-ID]/ likers?ref=about`

Our experiments have shown that the download of fanpages' users needs 2m30s up to 3m30s (depends on Internet and browser speed), where up to 1.213 fans can be achieved.

3.3 Crawling of private Likes

Lists of fanpages' fans are not complete due to privacy reasons. To improve the dataset, reactions of posts from fanpages can be used. According to [GO15], only 0.11% of users create a reaction to a post. With the following URL, the reactions of posts can be crawled:
`https://mbasic.facebook.com/ufi/reaction/profile/browser/?ft_ent_identifier=[Post-ID]`

The post-id can be extracted by the Facebook-API. The page lists the reactions of 10 users. By clicking on the link "show more", next 10 users will be listed. The parameters in the URL can be changed from "`limit=10`" to "`limit=5000`"[4]. This leads to the listing of much more reactions. An analysis has shown that reactions for posts can also be used to guess which fanpages were liked by people that marked their profile as private. Therefore, a sample of 120 reactions had been tested manually to answer the question whether people that interact with posts of fanpages also like them. The result shows that 87.7% also liked the fanpages if they liked a post. 12.3% did not like the fanpage.[5] This result can be transferred to private profiles, so that 87.7% might like the

[4] The maximum amount can change over time. This has to be tested.

[5] Used sample with post-id: 10155544812097244

fanpage if they liked a post. We confronted Facebook with the problem, that a simple user can download up to 900.000 reactions per hour. It has not been fixed yet, the security team of Facebook evaluated the approach as not being important for privacy.

3.4 Browser-History-Stealing and derive identity

For a successful attack, an attacker has to find out which fanpages a user has visited before. To do this, he can use the browser-history-stealing-attack. This is a side channel attack where an attacker can extract previous visited websites by using indirect access to the browsers history [RU14]. The technical background for this attack is the possibility to colorize letters depending on previous visited links. Coloring can be done with the CSS-formatting "`a:visited {color:blue;}`". An attacker could use this information to decide whether an URL, in this scenario a fanpage, has been visited before. The complete approach to create a fake-captcha can be found in [RU17].

If an attacker combines the knowledge of visited fanpages with the previously crawled dataset of users' Likes, he or she will be able to derive the identity of a person. Therefore, the attacker has to prepare a website with the fake-captcha, which has to be solved by visitors. The attacker assumes that (1) the visitor solves the captcha correctly, (2) he/she has a Facebook profile, (3) he/she uses Facebook with the current browser, (4) he/she used the Like button for fanpages before, (5) he/she visited the fanpages, (6) the browser stored visited URLs and (7) the attacker has already crawled Likes of the user so that his/her Likes could be found in the attackers' dataset. If (1) to (7) are fulfilled, the attacker can find a possible Facebook profile that might belong to the person. To do that, the attacker needs several rounds of captchas to check which fanpages might be visited before.

4 Results

During the research, we collected 7.092M Likes of 595,777 fanpages using two different methods. First, we directly crawled 350,000 profiles and their lists of Likes. The starting points to crawl profiles were 537 public groups whose header was in German language. Then the friends of the members were also crawled, until 350,000 profiles were achieved. This leads to a crawled popularity of ~300,000 German profiles and ~50,000 profiles, which are not located in Germany. Additionally, we collected Likes through the crawling of fanpages and their corresponding Likes. Finally, we collected 7.092M Likes, distributed over 927,803 people.

Tab. 1 shows the amount of users that can be identified using their Likes of fanpages. The higher the minimal known amount of Likes per user is, the higher will be the amount of identifiable users. If we consider that users made at least 4 Likes in the past, over 99.91% can be identified using their Like fingerprint. In general, every German user has left over 28 Likes on average [SO13]. Next to this, the general user creates two new Likes every month [SA11]. With the assumption of the amount of 28 Likes, 99.999% can be identified using their fingerprints.

Due to the direct crawling of fanpages and their Likes, each crawling process collects one Like per corresponding user. From them, no complete list of fanpages had been observed. This leads to a dataset that contains lots of users where only one Like of a user is known. From a statistical point of view, the result where only one Like per user exists ($X = 1$) is not relevant. Finally, the result of $X >= 1$ does not reflect the popularity because of missing values.

Minimal amount X of known Likes per user	**Amount of identifiable profiles using X Likes**	**Absolute amount of profiles with at least X Likes**
X >= 2	95.79%	665,488
X >= 3	99.81%	363,517
X >= 4	99.91%	205,463

Tab. 1: Fingerprint with Likes

For the experiments, a fake-captcha had been implemented to show the feasibility of the browser-history-stealing-attack and to show that it is possible to get the information whether someone likes a fanpage or not. It is also possible to use automatically approaches to get the information of previous visited fanpages [RU14].

5 Protection

5.1 Protection by social networks

The crawling experiment has shown that an attacker has the possibility to determine visited fanpages where users with a private profile interacted with. That is problematic because users might assume that their interactions also were private. The limitation of accessible profiles while the social network detects an unusual behavior, is one possibility to protect users' data. Another approach would be the limitation of users' private data by default. If users have to decide which data they want to publish and everything else stays private by default or is only accessible for friends, this would be a better barrier against an attacker.

At Facebook, a crawler can access up to 10,000 profiles in a period of 12 hours. This means that a user visits up to 14 profiles every minute during 12 hours. An attackers profile could be suspended much faster because the visit of 10,000 profiles cannot be a usual behavior. The time of the hold of 7 days could be increased.

5.2 Protection by browsers

The concept of a fake-captcha still works in every current browser, where a browser history will be stored. The problem is that the distinction between not-visited and visited websites using colors is a very basic concept of browsers. As long as there is an existing distinction, it can be used by an attacker to capture the information of previously visited websites. Users could disable the distinction between previous visited URLs. For instance at Mozilla Firefox, the setting in "`about:config`" of the parameter "`layout.css.visited_links_enabled`" could be set to false. This disables the distinction for all URLs. The disadvantage of this method is the deactivation of the distinction of all URLs, not only social media URLs. The attack can be fixed by a limitation of the distinction of previous visited social media URLs. If browsers do not make a distinction between previous visited social media websites, the combination of browser-history-stealing and the attacker's dataset will not work. A browser extension could do this as well.

5.3 Protection by users

"[Participants] may not see the information they provide as a threat to their future at present" [WE11]. Only 40% use the possibility to make their profile private "and few users change the default settings" [KO13]. One consequent protection would be not to publish private data in social media profiles. This can be realized by avoiding social media. But that is not a feasible solution for everyone. Before publishing data, users should understand how the privacy mechanisms of the used social network work. The first step is the maximization of the privacy by setting each kind of data to be private. While using social media, users should leave as less data marks as possible. Apart from the usernames and images, no other private data like mobile numbers, emails or dates of birth should be published for everyone.

Next to this, the visibility in search engines should be disabled. With this setup, search engines do not get a possible access to the private profile. Independently from the privacy settings at Facebook, the names of users and their user images are generally visible for other participants of the network. That should be changed by default. Another clue is the accepted amount of friends at a social network. "The category 'friend' is very broad and ambiguous in the online world; it may include anyone from an intimate friend to a casual acquaintance or a complete stranger of whom only their online identity is known" [DE09]. According to Rosenbury et al. [JU08], over 30% of users are willing to accept friendship requests from foreigners. At Facebook, a profile has at least access to his friends and to their friends. Users should always know that the visibility of data that is explicitly marked to be for friends, is available for a much wider amount of users than assumed.

7 Acknowledgments

This work has been supported by the Weizenbaum Institute for the networked society. We want to thank Jens-Peter Redlich for his support during the research. Additionally, the third author had been supported by the China Scholarship Council (Grant No. 201706775022).

Literature

[LI11] Y. Liu, K. P. Gummadi, B. Krishnamurthy and A. Mislove, "Analyzing facebook privacy settings: user expectations vs. reality," in IMC '11 Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference , Berlin, Germany, ACM, 2011, pp. 61-70.

[WE11] Z. Weinberg, E. Y. Chen, P. R. Jayaraman and C. Jackson, "I Still Know What You Visited Last Summer," in 2011 IEEE Symposium on Security and Privacy, Mellon University, 2011, pp. 147-161.

[RU14] S. Rüdian, "Browser-History-Stealing Ein Angriff auf die Privatsphäre," Humboldt University of Berlin, Berlin, 2014.

[SH08] Shmatikov, Arvind, Narayanan and Vitaly, "De-anonymizing Social Networks," in In Security and Privacy, University of Texas, 2009, pp. 173-187.

[NA08] A. Narayanan and V. Shmatikov, "Robust De-anonymization of Large Datasets (How to Break Anonymity of the Netflix Prize Dataset)," in IEEE Symposium on Security and Privacy (sp 2008), Texas, 2008.

[JI14] S. Ji, W. Li, M. Srivatsa and R. Beyah, "Structural Data De-anonymization: Quantification, Practice, and Implications," in Conference on Computer and Communications Security, Scottsdale, Arizona, USA, ACM SIGSAC, 2014, pp. 1040-1053.

[YO15] W. Youyou, M. Kosinski and D. Stillwell, "Computer-based personality judgments are more accurate than those made by humans," Department of Psychology, University of Cambridge, Cambridge CB2 3EB, United Kingdom; and Department of Computer Science, Stanford University, Stanford, CA 94305, Proc Natl Acad Sci USA, 2015, pp. 1036-1040.

[SO13] Socialbakers, "Cutting Through the Crowds on Facebook News Feeds," 09 04 2013. [Online]. Available: https://www.socialbakers.com/blog/1561-cutting-through-the-crowds-on-facebook-news-feeds. [Accessed 23 05 2017].

[SA11] M. Saleem, "Visualizing 6 Years of Facebook [INFOGRAPHIC]," 13 09 2011. [Online]. Available: http://mashable.com/2010/02/10/facebook-growth-infographic/. [Accessed 23 03 2017].[FA14] Facebook, "Common Scenarios for using the

Graph API," 14 07 2014. [Online]. Available: https://web.archive.org/web/20140714123747/https://developers.facebook.com/docs/graph-api/common-scenarios. [Accessed 12 06 2017].

[SE17] Seorld, "Facebook crawlen" 01 05 2017. [Online]. Available: https://seorld.com/blog/social-media/facebook [Accessed 12 06 2017].

[GO15] J. Gottke, "Infographic – Average Facebook Page Performance February 2015," 19 03 2015. [Online]. Available: https://www.quintly.com/blog/2015/03/infographic-average-facebook-page-performance-february-2015/. [Accessed 12 06 2017].

[RU17] S. Rüdian, "Deanonymisierung durch soziale Netzwerke am Beispiel Facebook," Humboldt University of Berlin, Berlin, 2017.

[KO13] M. Kosinski, D. Stillwell and T. Graepel, "Private traits and attributes are predictable from digital records of human behavior," Free School Lane, The Psychometrics Centre, University of Cambridge, Cambridge CB2 3RQ United Kingdom; and Microsoft Research, Cambridge CB1 2FB, United Kingdom, Proc Natl Acad Sci USA, 2013, pp. 5802-5805.

[DE09] B. Debatin, J. P. Lovejoy, A.-K. Horn and B. N. Hughes, "Facebook and Online Privacy: Attitudes, Behaviors, and Unintended Consequences," in Journal of Computer-Mediated Communication, Blackwell Publishing Ltd, 2009, pp. 83-108.

[JU08] K. Jump, "A new kind of fame," 21 07 2008. [Online]. Available: http://www.columbiamissourian.com/news/local/a-new-kind-of-fame/article_d5a1e536-bdf9-57b3-b5ad-49a28f381dab.html. [Accessed 10 03 2017].