

Uncovering Hidden Social Relationships Through Location-Based Services: The Happn Case Study

Adriano Di Luzio
Computer Science Department
Sapienza University of Rome
Email: diluzio@di.uniroma1.it

Alessandro Mei
Computer Science Department
Sapienza University of Rome
Email: mei@di.uniroma1.it

Julinda Stefa
Computer Science Department
Sapienza University of Rome
Email: stef@di.uniroma1.it

Abstract—Every day we publish an impressive amount of information about ourselves online: Pictures, videos, tweets, posts, and check-ins. These (big) data tell the story of our public lives and, at first, are seemingly innocuous. Nonetheless, they reveal a great deal of information about the things that we wished to remain private too. The goal of this work is to demonstrate how an attacker can exploit these data to jeopardize our privacy. In particular, we focus on Happn, a location-based dating app that counts millions of users worldwide.

Happn's goal is to let us find other users in our surroundings. It shows us their first name, age, and gender, their pictures, and their short bio. In addition, it tells us their approximate distance from our current location. By exploiting this information we show how an attacker (a stalker, a data broker, an advertisement company, etc.) can expose the private lives of Happn users and pinpoint their geographic position accurately and in real-time. In fact, we demonstrate how he can get even more detailed knowledge about the users by targeting all the population from a specific region (a city or a country) to discover their daily routines, where they work, where they live, and where they have fun. Finally, we show how to find out the social relationships between the Happn users: Uncovering who their friends, colleagues, relatives, and flatmates are.

I. INTRODUCTION

At the end of 2017 the Internet counted around 4 billion users [3] — more than half of the world's population. Every day, these users produce quintillion bytes of new data: Pictures, videos, Tweets, Facebook posts, Foursquare check-ins, and Tripadvisor reviews. Most of these data are available online and are publicly accessible. As such, they create huge security and privacy implications that are seldom mentioned. The goal of this work is to shed light on some of these implications, by showing how an attacker can exploit these data to jeopardize the privacy of millions of individuals. In particular, we focus on location-based services.

Location-based services play a primary role in our daily lives. Through Whisper, Foursquare, and Tinder we share our secret thoughts, we look for the recommended places to eat, and we seek our mate for life. While doing so we disclose a great deal of information about ourselves: Who we are and what we like, where we hangout with friends, and what is our favourite pizza place. But, sometimes, we reveal a lot more than what we originally meant to.

Happn, for example, is a location-based dating app that counts more than 37 million users worldwide [2]. Its goal is simple: To let users find the others that cross their path —

someone they noticed in a university class, on a train, or at a concert. For each user in our surroundings Happn shows us their first name, age, job, and a short bio. In addition, it displays their profile pictures and their current distance from us.

In this work we demonstrate how this seemingly innocuous information reveals a great deal of the private lives of the Happn users. In Section III we show how an attacker (*e.g.* a data broker or an advertisement company) can harvest the personal details of millions of worldwide Happn users from the comfort of his house — exploiting Happn's big data to jeopardize their privacy. Then, we show how he can trilaterate their geographic position, accurately and in a handful of seconds. In Section IV we demonstrate how an attacker can target all the users from a wide geographic area (*e.g.* a city, a district, or a country). To do so, we focus on 10 thousand Happn users from Rome, Italy, for 5 consecutive days. As a result we discover their daily routines and where they live, work, and hangout with their friends. Worse, we uncover their social relations with the others as well.

II. RELATED WORKS

The privacy-related issues that arise with location-based services attracted the interest of several researchers. Most services, indeed, show to their users the distance between their selves and the others. An attacker can repeatedly acquire this information from different locations to *trilaterate* the position of a target. Li *et al.* [8], as an example, reliably geo-locate 30 volunteers using WeChat, Skout, and Momo. Wang *et al.* [13] study Whisper, an anonymous network where users share their thoughts with the others nearby. They pinpoint the location of the *whisperers* with an accuracy of roughly 0.3 km. Ding *et al.* [6] analyze the one-week mobility traces of the WeChat users in mid-town Manhattan to identify their mobility patterns. Chen *et al.* [5] work on users characterization, analyzing on a large scale the activity of Momo users. Other, more formal approaches have also been presented. Peng *et al.* [10] rely on number theory to design two localization attacks that, potentially, can target any location-based social network. Polakis *et al.* [11] formalize the adversary model and audit various location-based services to evaluate their security.

In this work we take several steps further. First, we present a generic de-anonymization methodology based on trilateration:

It iteratively pinpoints the geographic position of any target in real-time and with an accuracy comparable to the GPS under open sky; it can be potentially used against any location-based service that reports the distance to the other users; it has proved to perform well even with a constant number of distance reports and, as such, it is scalable. Then, we build upon this methodology: We exploit its scalability to target thousands of Happn users, all at once. We pinpoint 10 000 users from Rome, Italy once per hour for 5 consecutive days and we study their mobility patterns and their routines. Then, we uncover where they live, where they work, their favourite places around the city, and their social connections as well.

III. HAPPN & USERS PRIVACY

Happn is a popular location-based dating app available for Android, iOS, and Windows Phone. It lets people find the others that cross their paths: Someone they noticed on the bus, while commuting to work, or in their favourite pub. It was launched in Paris in January 2014 [1], and 4 years later it counts more than 37 million active users worldwide [2].

Happn requires its users to register through a Facebook account. After setting up their dating preferences (gender, minimum and maximum age, etc.), the app welcomes its users into their *timeline*: A list of other users that are currently nearby, along with their bio, pictures, age, and profession. From the timeline, a user can take a look at the profiles of the others (*e.g.* their current distance, the last time they used the app, the songs they listen to, etc.) and can express interest in them. If the interest is mutual Happn notifies both users and lets them communicate through instant messaging and short voice mails.

The profile of any Happn user contains, in particular, her current distance from us. For privacy reasons, the app does not tell use the exact distance. Instead, it adopts a systems of concentric circles of increasing radii to indicate that someone is within 250 m, 500 m, 1000 m, 2000 m, etc. A user 300 m away would be placed in the 500 m circle. Another, 70 m away, would be reported as within 250 m (the closest that any user is reported).

The goal of this study is to investigate to what extent an attacker can compromise the privacy of Happn users (and possibly of other location-based services too). For this reason, we play the role of someone interested in discovering as much as possible about them: *E.g.* an advertisement company that wants to profile mid-twenty users from Rome, New York, or Berlin; a governmental organization that wants to uncover the hidden social relationships between the people that threaten their country; a stalker that wants to track the behaviour and the whereabouts of one or more individuals. We assume that the attacker controls one or more Happn users, that he created by registering to the service with as many Facebook accounts. We do not require him to be physically close to his targets, neither to have any a-priori knowledge about them.

Under the hood the app communicates with its backend servers through HTTP over TLS. To perform our study, we reverse-engineered the communication protocol by putting

an HTTPS proxy between the Android app and the Happn backends (performing a man-in-the-middle attack). Then, we built a customized Happn client that mimics the mobile application. We used it to programmatically set our position, to retrieve the list of other users around us, and to collect their personal information.


A. Profiling Happn Users

We begin our analysis on the privacy of Happn users by aiming at profiling them. Once in a while, the Happn mobile application sends our current position to its backend servers. Then, it requests a new list of *crossings* to update our timeline with the users that are currently nearby. For each user, the backend communicates to the app her first name, age, and gender; her profession and her workplace; her profile picture and the other photos that she decided to share (possibly, her Instagram profile); the last time she used the app and her current distance from us. The application displays these details through its user interface. An attacker can harvest them by using one of many alternatives: Eavesdropping on the communication between his mobile application and the Happn servers; scraping the particulars off the application interface by exploiting an Optical Character Recognition (OCR) software; directly requesting them to the backend with a customized Happn client, as we did. As a result, the attacker can thus build a dossier of nearby Happn users containing personal information, pictures, and the neighbourhood where they crossed his path. Besides, he can also infer their dating preferences. Each user appears in someone else's *crossings* list only if their dating preferences match (in terms of age, gender, etc.). So, if a man in his thirties appears in the timeline of a woman of his age, then we know that they expressed interest in their reciprocal genders, for a range of ages that include the thirties.


In addition, the backend also communicates to the app something that the interface does not display: The Happn identifier of each user nearby, and, more importantly, their Facebook ID. Through Facebook, an attacker can further extend his dossiers of the Happn users by including their last name, their Facebook profile picture, and all the other information that the user chose to share with the public. Sometimes [7], [4], this information includes their list of friends, the pages they like, the movies they watch, their interests, etc. It is unclear why Happn sends the Facebook identifiers to its mobile applications. By decompiling the Android version of the app, in fact, we discovered that this field is ignored by the code that processes the users' profiles. Still, these identifiers pose at great risk the privacy of Happn users, revealing information that would otherwise be difficult to uncover. They expose Happn users to a greater risk to be profiled or worse stalked by an attacker.

B. Geographic De-Anonymization of Happn Users


We also investigate whether an attacker can geographically de-anonymize the Happn users, *e.g.* by discovering their whereabouts or where they live and work. To do so, we exploit the distance information that Happn shows in the profiles of



(a) The initial trilateration of the target T from A , B , and C .



(b) D is the centroid of the polygon from Figure 1a. The new t-point from D roughly doubles the accuracy of the result.



(c) The t-points north, east, south and west of D further improve the result. For clarity we only show E , south of D .

Figure 1: An iterative trilateration methodology to geographically de-anonymize Happn users.

its users. As we have seen in Section III, this information places nearby users within circles of increasing radii (250 m, 500 m, 1000 m, etc.) centered at our current position. This system aims at protecting their privacy, and reveals a rough approximation of their position — in the best case within an area of $250^2 \cdot \pi = 196\,250\text{ m}^2$. Nonetheless, we show how an attacker can pinpoint the position of any target with an accuracy of 6 m.

Towards our goal we use an iterative methodology based on *trilateration* (Figure 1). The process begins by acquiring the distance to the target from multiple locations, the trilateration points (t-points for short). By considering these points and the reported distances, the trilateration creates a polygon containing the target (Figure 1a). The accuracy of the result depends on the position of the target and on the t-points. Nonetheless, an attacker can further improve upon it: First he selects the centroid of the polygon containing the target; then, he asks for the distance to the target from that point (Figure 1b); finally, he repeats the previous steps from four additional t-points located north, east, south, and west of the polygon’s centroid (Figure 1c). This process iteratively reduces the area of the polygon containing the target and can be repeated until it meets the desired accuracy.

To exploit this methodology an attacker needs to request the distance to the target from each trilateration point. This, in turn, requires a different user for each request. Happn forbids, indeed, its users from moving too fast to different places. For this reason, the attacker cannot relocate a single user to the

Table I: The trilateration accuracy after each iteration.

Iteration	Distance Reports	Error (m)
0	3	79.66
1	8	25.27
2	13	6.21

different t-points. Even so, he can use multiple *decoy* users (*i.e.* multiple Happn accounts) and *spoof* their position — in our case, we programmatically place them at the t-points through our custom Happn client.

Generally speaking, the accuracy of the geo-localization increases with the number of distance reports. Nonetheless, our methodology has proved to be effective even with a small number of t-points (thus requiring a small number of accounts). As an example, we tested its real-world performance by geo-localizing a volunteer that used the application on a real Android device from a fixed location in Rome, Italy. Table I shows the results of the experiment. The initial trilateration from three locations chosen at random within Rome results in a distance between the polygon’s centroid and the real-world position of the user of 79.66 m. The succeeding iterations progressively increase the accuracy of the result, pinpointing the target to an area roughly the size of a basketball court with an error of 6.21 m — in comparison, the accuracy of the GPS of a smartphone under open sky is 4.9 m [12].

The entire process, overall, lasts a handful of seconds — the time to request the distance reports to Happn. It requires an account for each request and the identifier of the target. This identifier can be discovered from the attacker’s crossings or by enumeration: Happn assigns increasing numeric identifiers to its users and the attacker can ask for the details of any user, not just for those in his surroundings. Thus, with this methodology we show how an attacker can discover the accurate location of any Happn user regardless of their gender, their age, their dating preferences, and of whether they previously crossed his path.

IV. UNCOVERING HAPPN SOCIAL RELATIONSHIPS

Location-based services jeopardize our privacy in ways that we hardly expect. In the previous sections we showed, indeed, how an attacker can collect a detailed profile about any Happn user (including their personal information, their interests, their social media profiles, etc.) and also how he can instantaneously pinpoint their accurate whereabouts. This information, in turn, leads to more concerning results. In the remainder of this work, we show how an attacker can target *en masse* all the Happn users from a wide geographic region (*e.g.* an entire country). In particular, we track, profile, and then pinpoint almost 10 000 Happn users from Rome, Italy. We uncover their daily routines, their home and work addresses, where they like to hangout on weekends, and who their friends, their relatives, and their colleagues are. Through Happn we not only expose their private lives, but their social connections as well.




Figure 2: The genders and ages of the Happn users from Rome.

A. Experimental Setup

To carry out this task, we collected the profiles of thousands of Happn users from Rome, the Italian capital. We did so by creating 20 Happn accounts and then programmatically setting their position at as many popular places within the city: University campuses, train stations, places where young people usually hangout, well-known bars, pubs, restaurants, etc. From each location, we repeatedly acquired the list of nearby users. To maximize the number of results, we made sure to shuffle the positions of our accounts and to diversify their gender and dating preferences. After a few days we counted 9215 users — all those who either live in Rome, or study, work, or spend a significant portion of their time there. For each of them we stored their identifiers, first name, age, and workplace. In Figure 2 we present a brief analysis of their demographic features, investigating the distributions of their age by gender. We note that the number of men is roughly 7 times larger than that of women (8006 against 1209). Both the trends present their maximum around the age of thirty, starting from 18 (the minimum age required by Happn) and ending around 65 years.

After acquiring their profiles, we continuously pinpointed the positions of the Happn users from Rome. To do so, we programmatically arranged our 20 users around the city, at 20 fixed trilateration points. Then, we used our trilateration methodology to recurrently uncover the exact location of each target. We repeated this task, once per hour, from 2017-04-07 20:00:00 to 2017-04-12 19:00:00 (*i.e.* for 5 consecutive days). As we have seen in Section III-B the trilateration starts from some initial locations. Then, it improves upon its accuracy by requesting additional distance reports to Happn. It requires, however, a new account for each request: After we place one of our decoy users to a trilateration point, we can't move it to a new position unless we wait for a cool-off period. Since we were monitoring a high number of individuals, we limited the trilateration to its initial iteration and we fixed our 20 users at their starting locations. Nevertheless, this methodology has proven to be quite effective.

Of the 9215 Happn users that we observed, we focused on the 6373 that use the application continuously (*i.e.* by interacting with it multiple times per day). The trilateration placed almost 1000 of them within the area of a square of side 30 m (*i.e.* less than the surface of an Olympic-size swimming pool). The 51% of the users have been located, at best, within a square of side 70 m (roughly the area of an American football field). More than 82% of the targets have been pinpointed in

a square of side 140 m, the average area of a Manhattan city block.

The geographic de-anonymization has thus proven to work effectively also on a city-wide scale. We actively tracked for 5 days the whereabouts of thousands of Happn users around Rome. For this task we did not require any previous knowledge on the users except for their Happn identifiers. In addition, we only exploited a small, constant number of decoy accounts. For the majority of targets, we uncovered their position with an accuracy up to the city block — enough for an attacker (*e.g.* a stalker) that aims at monitoring or harassing one or more individuals. Roughly half of the times, we discovered the particular building where they were currently staying at. In some cases, we pinpointed them within the area of an office, or of a small house.


In turn, this leads us to our next observation. An attacker that focuses on the pinpointed locations of a single user can also uncover her routine, her favourite places, and where she is more likely to be at any time of the day. So, the attacker can find out where she spends the night (*i.e.* where she lives), where she has lunch, where she goes every morning (*i.e.* where she works and what is her daily commute), and where she hangouts on weekends. Furthermore, an attacker can leverage this information to discover who she hangouts with, identifying her colleagues, her flatmates, her friends, and her relatives — uncovering Happn's underlying social network.

B. Happn's Underlying Social Network

Let us say that Alice and Bob, two Happn users, have been placed by the trilateration methodology within the same geographic area, at the same time. Then, it is possible that Alice and Bob know each other. The smaller the area and the more often this happens, the more confident we are about their relation. An attacker can exploit this information to uncover the social relation between Alice and Bob, extending what he already knows about them.

In this study we targeted the Happn users from Rome, trilaterating their location once per hour. While doing so, we not only discovered their whereabouts and their daily routines, but we also built a social graph of their personal relationships with friends, relatives, and colleagues. The nodes of the graph represent Happn users and the edges between them reflect their relationships. There is an edge between two users if we pinpointed them together in an area smaller than a Manhattan city block at least once. In addition, edges are labeled and weighted. The weight counts the number of times the users were pinpointed together. The labels, instead, reflect the categories (friends, colleagues, or flatmates) of their social relation, and are based on where, when and how often the users have met. As an example, users that exclusively meet in business hours in a office tower are likely to be colleagues. Similarly, users that meet at deep night in an apartment building are probably flatmates.

The resulting social graph consists of 5583 vertices and 14 497 edges. Its diameter is 13, its average path length is 2.06, and its clustering coefficient is 0.193. These measurements



(a) Log-log ICDF of degrees. (b) Log-log ICDF of weights.

Figure 3: Log-log ICDFs of node degrees and edge weights.

are compatible with the structure of an online social network [9]. In Figure 3 we depict the log-log inverse cumulative distribution functions of the nodes degrees and of the edges weights. The overall trends are, as one would expect, both decreasing. The distribution of Figure 3a is consistent with the power-law of degrees in a social network [9]. There, we note that 979 nodes have a degree of 0 (*i.e.* are disconnected from the others). They correspond to users that have never been pinpointed together with anyone else. By inspecting their trilateration results, we discovered that they live outside of Rome and that they, in fact, sporadically frequent the city. In Figure 3b instead, we note that the minimum weight is 1, denoting that users in the graph share an edge only if they have met at least once.

This social graph reveals plenty of information about its users. An attacker can exploit it to uncover the (possibly hidden) social relationships of single individuals or to target entire sub-communities of users. As an example, he could focus on those that profess the same religion to discover whether they meet outside of their place of worship; he could aim at those that belong to the same political party, to threaten, stalk or influence them; he could pinpoint the positions of police and military forces around a city; he could select IT workers, lawyers, or physicians to investigate whether they have contacts with competing companies.

While analyzing this network we found many interesting case studies. In the remainder of this section we tell some of their stories and we demonstrate the extent to which an attacker can jeopardize the private lives of his targets. We modify the names of the interested users to preserve their privacy. However, we keep their ages, their jobs, and their social connections intact. In addition, we note that an attacker can also leverage the Facebook profiles of all users, as provided by Happn. They include without fail their last names. Depending on their privacy settings, also their entire list of friends, of interests, and of previous posts.

Paola Paola is 26 and we know from her Happn profile that she works in financial services. Her Facebook profile is mostly private, but it tells us that she is originally from Tuscany. The trilateration reveals the block where she lives in and her friends. Anna, for example, is 26 too. On a Saturday night, Paola slept-over at Anna's house. They spent their Sunday there, together with Claudio (Anna's flatmate, lawyer, 31). On Saturday night he arrived home at 5 a.m. Earlier, he was in a club in the city center with Giuseppe, a friend of his. We also

know Angela (32) and Mario (42), two of his colleagues, that exclusively meet him in business hours of working days.

Sandro Sandro is 24. His Happn profile does not disclose his job, however his Facebook profile reveals that he is a university student. We know from his trilateration that he is studying biomedical engineering. He spends, indeed, his business days at the Institute of Biomedical Technologies of Rome. The institute is also conveniently close to his home, where he spends most of his free time. We know the exact building of his apartment, located in a neighbourhood popular for university students. He shares it with Antonio (33), Gaia (26), and Michele (19). Antonio and Gaia work in the university hospital, respectively as a computer technician and as a specializing physician. Michele, instead, is a university chemistry student.

Giulio Giulio is 24 years old and he works for a travel agency. He shares a flat with Maria (28) and Stefano (29), and we know its location within a few meters. On business days, he works at the agency, of which we know the address. In the evening, he likes to have a drink with his friends in a popular Roman bar.

V. COUNTERMEASURES

Protecting the privacy of location-based services users is a challenging task. Happn, for example, protects its users in two ways: On one side, it limits the amount of personal information displayed (*e.g.* their first name only) to make their identification harder. On the other, it reports their distance after rounding it to multiples of 250 m (*i.e.* it reports them within circles of increasing radii). Nonetheless, an attacker can discover a concerning amount of information about any Happn user and, worse, can target an entire geographic region (*i.e.* a nation) to massively monitor thousands of individuals. In the remainder of this section we propose several countermeasures to address these issues.

Trilateration An attacker that wants to pinpoint a user needs to acquire several distance reports in a short time. To do so, he requests the distance to the target by using multiple Happn accounts placed in as many different spots around the city. A possible countermeasure to this trilateration attack could be to detect the repeated distance requests to the target profile (regardless of the source account) and to block them or, alternatively, to add a random amount of noise to the reported values.

Registration New users register to Happn through their Facebook account. However, creating fake Facebook accounts for this purpose has proved to be trivial. A possible countermeasure to prevent the registration of fake users could be to add a verification phase to the registration process: *E.g.* by requiring a unique phone number or by validating their details and their pictures.

Rate-limiting While monitoring the Happn users from Rome we performed almost 10 000 requests per hour from each of our 20 accounts. The entire process completed in a short time

and we executed numerous requests per second from each account. A possible countermeasure could be to automatically limit the rate by which their backend endpoints can be requested (e.g. distance reports, user profiles, etc.) both in terms of IP addresses and Happn accounts.

Users profiles An attacker can request the profile of any Happn user, regardless of whether she previously crossed his path. A possible countermeasure to prevent this behaviour could be to only display the profiles of users that have been in his surroundings before. This would also forbid user enumeration.

Least privilege The profile of each Happn user contains a reference to her Facebook, Twitter, Instagram, and Spotify profiles. The Facebook identifier, in particular, provides invaluable information to an attacker that aims at jeopardizing the privacy of a user. To the best of our knowledge, the Happn mobile apps ignore these identifiers when processing the information of the user profiles. As such, a possible countermeasures could be to avoid sending them back to the mobile clients.

VI. ETHICAL CONSIDERATIONS

In this study we analyzed Happn and the privacy of its users. While doing so, we collected the personal details of almost 10 000 individuals from Rome, Italy. Gathering this information without the users' consent is a potential ethical issue. For this reason, we protected with particular care the privacy of the individuals that are part of this research: We only collected the minimum amount of information required to perform this study. We stored all the data that we collected in an encrypted database. We did not share them with any third party and we will destroy them as soon as this work is finished. In addition, we emphasize that the Happn users will benefit from this research: Our goal is to shed light on some of the important issues related to their privacy, both online and in real-life.

VII. CONCLUSIONS

This work demonstrated to what extent the seemingly innocuous information that we reveal about ourselves online put our privacy at great risk. We focused on location-based services and social networks and we showed on how an attacker can exploit their impressive amount of information (i.e. their big data) to jeopardize the privacy of millions of users. Happn, in particular, is a popular location-based dating app that counts millions of users active every day, from all over the world. We started off by showing how an attacker can harvest the personal information of any Happn user, regardless of their gender, their dating preferences, and of their geographic position. Then, we exploited Happn's distance reports to accurately pinpoint the position of any target, in real-time. We showed how an attacker can target *en masse* all the Happn users from a sub-community (e.g. based on religion or politics) or from a wide geographic region for continuous periods of time. We focused on the Happn users from Rome, Italy — monitoring almost 10 000 individuals for 5 consecutive days. As a result, we uncovered their daily routines, where they work, where they study and

where they like to have a drink. Finally, we exposed their social relations, pointing out their friends, their colleagues and their flatmates.

Some location-based services, in fact, jeopardize the privacy of their users to a severe extent. We thus hope that the consequences of this work will be manifold. On one side, we want to improve the security of new and existing location-based services and social networks. On another, we aim at increasing the awareness of users' regarding both their real-world and their online privacy. Finally, we hope to show how the seemingly innocent personal details that we reveal can be exploited to draw a detailed, concerning picture of our private lives.

REFERENCES

- [1] *Happn - About us*, 2014 (accessed Dec 27, 2017). <https://www.happn.com/about>.
- [2] *Happn - Press release*, 2017 (accessed Dec 27, 2017). <https://www.happn.com/it/press>.
- [3] *World Internet Users and 2017 Population Stats*, 2017 (accessed Dec 27, 2017). <http://www.internetworldstats.com/stats.htm>.
- [4] A. Acquisti, L. Brandimarte, and G. Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.
- [5] T. Chen, M. A. Kaafar, and R. Boreli. The where and when of finding new friends: Analysis of a location-based social discovery network. In E. Kiciman, N. B. Ellison, B. Hogan, P. Resnick, and I. Soboroff, editors, *ICWSM*. The AAAI Press, 2013.
- [6] Y. Ding, S. T. Peddinti, and K. W. Ross. Stalking beijing from timbuktu: A generic measurement approach for exploiting location-based social discovery. In *Proceedings of the 4th ACM Workshop on Security and Privacy in Smartphones & Mobile Devices*, SPSM '14, pages 75–80, New York, NY, USA, 2014. ACM.
- [7] R. Gross and A. Acquisti. Information revelation and privacy in online social networks. In *Proceedings of the 2005 ACM Workshop on Privacy in the Electronic Society*, WPES '05, pages 71–80, New York, NY, USA, 2005. ACM.
- [8] M. Li, H. Zhu, Z. Gao, S. Chen, L. Yu, S. Hu, and K. Ren. All your location are belong to us: Breaking mobile social networks for automated user location tracking. In *Proceedings of the 15th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '14, pages 43–52, New York, NY, USA, 2014. ACM.
- [9] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*, IMC '07, pages 29–42, New York, NY, USA, 2007. ACM.
- [10] J. Peng, Y. Meng, M. Xue, X. Hei, and K. W. Ross. Attacks and defenses in location-based social networks: A heuristic number theory approach. In *2015 International Symposium on Security and Privacy in Social Networks and Big Data (SocialSec)*, pages 64–71, Nov 2015.
- [11] I. Polakis, G. Argyros, T. Petsios, S. Sivakorn, and A. D. Keromytis. Where's wally?: Precise user discovery attacks in location proximity services. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, pages 817–828, New York, NY, USA, 2015. ACM.
- [12] F. van Diggelen and P. Enge. The World's first GPS MOOC and Worldwide Laboratory using Smartphones. In *Proceedings of the 28th International Technical Meeting of The Satellite Division of the Institute of Navigation*, ION GNSS+ 2015, pages 361 – 369, Tampa, FL, USA, 2015.
- [13] G. Wang, B. Wang, T. Wang, A. Nika, H. Zheng, and B. Y. Zhao. Whispers in the dark: Analysis of an anonymous social network. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, IMC '14, pages 137–150, New York, NY, USA, 2014. ACM.