User Interface Design for Privacy Awareness in eHealth Technologies

Isabel Wagner, Ying He, Duska Rosenberg and Helge Janicke
Faculty of Technology
De Montfort University, Leicester, UK
{isabel.wagner, ying.he, duska.rosenberg, heljanic}@dmu.ac.uk

***Abstract*—In this paper we investigate privacy issues relating to Human Computer Interfaces for mobile eHealth technologies. We present the Inform–Alert–Mitigate (I-AM) cycle, a novel approach to address privacy concerns that are associated with the use of these technologies. The I-AM approach supports the responsible innovation of new technologies. We demonstrate the effectiveness of I-AM by applying it to examples taken from mobile applications relating to personal health. We discuss three classes of applications: a) fitness trackers b) personal wellbeing applications and c) medical applications, and evaluate the privacy exposure of their users using representative applications from these classes. The paper evaluates the current privacy enhancing features of these applications against the identified risks and demonstrates how the I-AM approach can be applied to yield additional and more effective privacy protection for these technologies.**

I. Introduction

Mobile eHealth applications such as fitness trackers, weight loss or personal wellbeing apps, and medical applications are widely used today and have unprecedented access to users' sensitive information. Although many users have privacy concerns, they may not actively consider privacy while using these apps. Privacy notices, required by data protection regulations [1], are used to help users make informed privacy decisions. They inform users about data collection, processing, and sharing and make all data practices transparent to the users. However, previous research suggests that privacy notices do not correspond to users concerns [2]. In addition, users struggle to understand privacy notices as these notices are usually complex, long, and difficult to follow.

Current privacy studies focus on analyzing an app's usage of sensitive resources. User perception and interpretation of privacy notices have not been well considered, for example, some privacy notices were found not to be grounded in user research [2]. The HCI community has realized the importance of user perception and has come to the consensus that modern interfaces for privacy sensitive applications must inform users about the sensitivity of the data and functions the current application is operating with [3], [4]. In the context of mobile apps, efforts have mainly been made to increase privacy awareness before an app is installed [5], and to improve how an app's required permissions are displayed [6].

Despite the previous work on the usage of sensitive resources and interface design for sensitive applications, designers have very little guidance to follow, because interface design has not been integrated properly into systems interaction design [7]. In this paper we argue that the integration of user interface design with the design of privacy sensitive applications has to take into account the characteristics of both the physical and the informational environments – the user's entire interaction space [8].

In the following, we present the Inform–Alert–Mitigate (I-AM) cycle, a novel approach that provides a user-centric approach to creating privacy-aware user interfaces for mobile users. To demonstrate the applicability of I-AM, we apply the approach to four selected eHealth applications: two fitness trackers, a weight loss app, and a medical application for diabetic users. We discuss potential privacy concerns for each of the apps, analyze how privacy-aware their interfaces currently are, and propose how they could be improved following the three stages of the I-AM cycle.

The remainder of this paper is structured as follows. After we review related work on privacy aspects in user interface design (Section II), we introduce the functionality of four selected eHealth applications for mobile phones in Section III. We go on to highlight privacy concerns both in general and as they manifest in the four apps (Section IV). In Section V, we introduce the Inform–Alert–Mitigate (I-AM) cycle, a novel concept for assessing and improving privacy aspects. We demonstrate the power of the I-AM cycle in Section VI by applying I-AM to the four selected eHealth apps, point out shortcomings, and identify improvements. Finally, we conclude in Section VII.

II. Related Work

a) Current state of privacy notices: Smartphone operating systems use privacy notices to help users make informed privacy decisions. Users see privacy notices as warnings during installation or at run time. Current research focuses on the user's privacy and security protection by leveraging application analysis such as recovering the application source code directly from its installation image [9], or apply security extensions that provide application specific privacy controls to users such as the application of permission-hungry applications in android platform [10]. These systems are useful for analyzing an app's usage of sensitive resources, however, they have not considered the user's perception and interpretation of the permissions [2]. Common user concerns concerning privacy notices are:

- The multitude of privacy notices make it difficult to process and recall [11].
- Some applications can access resources without providing a privacy notice [12].

- Users tend to ignore the privacy notice display as it sometimes appears after the users have decided to use that application [13].
- A series of semi-structured interview studies showed that users are unaware of privacy risks even after they read the privacy notices. They are not currently well prepared to make informed privacy decisions [14].
- Even users who pay attention to privacy notices have difficulty using them as sometimes the screens are jargon-filled, and the explanation of data collection purposes is either not stated or unclear [5].

This list clearly shows that there is a need to design better user interfaces to help users make informed privacy decisions.

b) Current state of privacy-aware user interface design: Schaub proposed a design space that considers both the audience's privacy notice requirements and system constraints [7]. He developed a taxonomy and notice approaches that help users understand the notice options available within a specific system context. Barkhuus studied contextually grounded reasons for users' privacy concern based on Nissenbaum's theory of contextual integrity [15]. Differently from previous work, Barkhuus [16] treated privacy as an information flow instead of static information sharing. Heng [17] has improved privacy notice dialogs by considering both control and awareness to address users' privacy concerns towards third-party apps on Facebook. He then adopted Privacy by ReDesign to study whether a user can adequately express their preferences for sharing and releasing information using the improved privacy notice dialog design.

There is a consensus in the HCI community that modern interfaces for security sensitive applications must inform users about the sensitivity of the data and functions the current application is operating with [3], [4]. Several frameworks and patterns have been developed [18], [19] to ensure that users understand and are presented with visual cues about the confidentiality considerations of data they are currently processing. In particular these concerns have been addressed in systems that employ Mandatory Access Control such as Multi-Level Security (MLS) [20]. Similar considerations apply for applications that handle sensitive data that may affect a user's privacy. In this paper we propose a comprehensive approach to address HCI design for privacy awareness and control in the context of eHealth applications for mobile devices. Prior work has proposed modifications to the Google Play store to increase privacy awareness before a user makes the decision to install an app. To this end, [5] propose to include a privacy summary in the app descriptions. Similarly, [6] propose a modification to the Google Play store that improves how an app's required permissions are displayed. Their suggestions are based on user expectations and surprise. Our approach focuses on the applications themselves and is therefore orthogonal to this prior work.

III. Selected eHealth Applications

To illustrate our approach for privacy-enhancing HCI design, we focus on three categories of popular eHealth applications: fitness trackers, weight loss apps, and medical applications. We use one or two representative apps from each category to illustrate their functionality (this section), describe privacy concerns arising from their usage (Section IV), and argue how the I-AM cycle can be used to improve privacy awareness and privacy protections (Section VI).

A. Fitness Trackers

Fitness trackers allow users to record data about their activities. The apps frequently rely on location information (e.g., GPS) to record distance and elevation changes; some also use external devices such as heart rate monitors. In addition to this core functionality, many fitness trackers provide maps of individual activities, statistics and graphs summarizing activities over a period of time, and comparisons with other users. We chose two representative apps in this category, Runtastic[1] and RunnerUp[2]. Runtastic is a popular fitness tracker that offers many features in addition to the core functionality, such as goals, challenges, and live tracking. To enable these features, activity data is stored on Runtastic servers. Runtastic is a commercial app owned by the adidas group. RunnerUp is an open source fitness tracker that offers a basic set of features. By default, RunnerUp stores activity data locally, but allows users to share activities with third-party websites, among them Runtastic. This means that users control which activities they share, when they do it, and with who. We will describe in Section VI how this concept can be extended to implement different kinds of privacy-enhancing features.

B. Weight Loss (Personal Wellbeing)

Weight loss apps are designed to help users achieve their dieting goals. They do this by providing food plans and allowing users to record food consumption, exercise, and weight. We chose the commercial app MyFitnessPal[3] as representative app for this category. MyFitnessPal visualizes a user's progress with charts and statistics, and offers a user community for support and motivation.

C. Medical Applications

Medical applications are intended to support users with a wide range of medical conditions, such as pregnancy, blood pressure, migraine, or diabetes. The specific kind of support varies with the medical condition. We chose Diabetes:M[4] as representative app in this category. The app offers diabetic users a tool to track their blood sugar values, carbohydrate intake, and insulin, and supports users with information about injection sites and an insulin calculator. The app can visualize variations over time, and provides reports that can help physicians make treatment decisions.

IV. Privacy Concerns

The privacy concerns for eHealth applications stem from the sensitivity of the data the applications handle, and depend on what data the apps are able to access, where the data is stored, how it is used, and how/with who it is shared. To assess the potential causes for privacy concerns for each app, we analyzed the descriptions on Google's Play store, the requested Android

[1] https://play.google.com/store/apps/details?id=com.runtastic.android

[2] https://play.google.com/store/apps/details?id=org.runnerup or https://f-droid.org/repository/browse/?fdfilter=fit&fdid=org.runnerup

[3] https://play.google.com/store/apps/details?id=com.myfitnesspal.android

[4] https://play.google.com/store/apps/details?id=com.mydiabetes

	Runtastic	RunnerUp	MyFitness-Pal	Diabetes:M
Collected data				
PII (name, email, address)	x	–	x	–
date of birth, age	–	–	x	–
height, weight, gender	–	–	x	x
activity data (time, duration, location, heart rate)	x	x	x	–
nutrition, caloric intake	–	–	x	x
medical data (blood sugar, blood pressure, etc.)	–	–	–	x
Android permissions				
location	x	x	x	x
device storage	x	x	x	x
camera, microphone	x	–	x	–
phone identity	x	–	x	–
network access	x	x	x	x
accounts	x	–	x	x
contacts	–	–	x	–
in-app purchases	x	–	x	–
Data use				
provide app functionality	x	–	x	–
in-house research	–	–	–	–
advertisers	x	–	x	x
sale	–	–	x	–
social media plug-ins	x	–	–	–
analytics, tracking	x	–	x	x
Data storage				
device	x	x	x	x
company servers	x	–	–	–
cloud servers	?	–	x	–
Data transfer				
using SSL/TLS	x	x	x	x
plaintext leaks	location	–	location, device info	–

TABLE I. PRIVACY CONCERNS

permissions, the privacy policies, and the traffic sent/received while using the app. We summarize our findings in Table I.

The types of data collected by apps – as far as acknowledged by the apps themselves – include personally identifiable information (PII) such as name and address as well as demographic data (e.g., date of birth, gender) and potentially sensitive data about activities (e.g., location, heart rate), nutrition, or medical data. It is interesting to note that there are differences in data collection even between apps in the same category (Runtastic and RunnerUp).

The Android permissions requested mostly correspond to the data that each app acknowledges to collect. However, there are some surprising permissions that do not obviously match collected data types. These include access to camera and microphone, access to the phone identity and accounts on the device, as well as access to the user's contact list.

The data collected by each app is used for multiple purposes. The usage of data to provide all or part of the app's functionality is expected and likely acceptable to most users. However, several apps also claim that they use non-personally identifiable information (PII) to give it to advertisers, sell it to interested parties, connect users to social media, or for analytics and tracking. These uses may be surprising for users and are less likely to be acceptable. The use of non-PII for these purposes does not provide much protection, because it has been shown that correlating non-PII with other data sources can lead to re-identification (see for example [21]).

Some apps store collected data locally on the device, while others transfer it to company or cloud servers. While the privacy policies usually claim that data is stored securely, the apps do not reveal more specific information and do not provide a way for users to verify this claim.

When data is transferred to the app provider or a third party, it is vulnerable to network-based eavesdroppers unless suitably encrypted. All apps encrypt the bulk data transfers; however, some of them leak plaintext information about a user's location or device.

Despite the extent of these privacy concerns, many users are neither aware of what data they give away and how it will be used, nor of how they could protect their privacy [2].

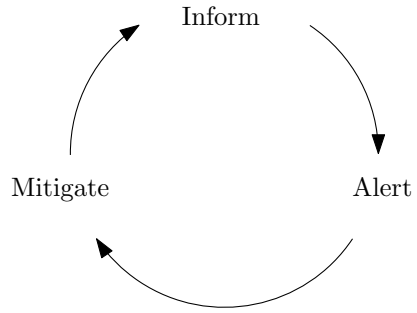

Fig. 1. Inform – Alert – Mitigate cycle

V. I-AM: INFORM – ALERT – MITIGATE

The Inform–Alert–Mitigate (I-AM) cycle provides a structured way to assess and improve how apps communicate privacy aspects and mitigation options to users. It can thus improve privacy awareness and give users effective control over their privacy. The I-AM cycle, shown in Figure 1, builds on existing best practices and frameworks [1], [5], [6], which all belong to the Inform stage. I-AM adds an Alert stage and a Mitigation stage in which the system will alert the user to current and ongoing privacy risks, and provide the user with mitigation options to minimize the incurred risk. The stages are designed to influence and build on each other. For example, an alert can be based on information about privacy policies and lead to mitigation actions being taken, which in turn will be used to inform the user of their efficacy.

A. Overview

c) Inform: The information step aims to inform users of potential privacy issues by making users aware of privacy policies and requested permissions. This is a typical part of the current design of privacy-aware Human Computer Interfaces and includes, for example, the display of requested permissions in the Google Play store [5], [6]. In I-AM, this is the starting point for designing privacy-aware interfaces for mobile eHealth applications. We consider the ability of a user to undertake her tasks as paramount and therefore propose a user-centric privacy approach that empowers users to make decisions about aspects of their privacy and effectively provide them with assistance and information to make privacy relevant decisions.

d) Alert: This step extends traditional approaches in that it uses contextual information about the current user environment and alerts the user about privacy risks that the user is exposing herself to. The contextual information could for example originate from the user's usage pattern, permission

usage by the app (such as accessing sensors or other sensitive information on the device), or network accesses. The alerts will flag up current issues that I-AM detects and allow the user to query the alerts to identify suitable mitigation tactics. The alerts may consist of visual cues, such as annotations to the app window or the device's status bar to indicate that sensitive information is being displayed and processed.

e) Mitigate: When queried about an alert, interfaces following the I-AM approach should provide the user with concrete actions that can mitigate the current risk-exposure. Some situations can also trigger automated mitigation actions undertaken by the system. For example an app transferring sensitive data over a network may indicate potential privacy problems. If such an event is combined with context information that the data transfer is not encrypted, it can trigger automatic responses (such as blocking the transfer). If a high-risk event is detected, but no automatic response can be provided, an I-AM interface should provide a set of mitigation actions for the user that are aimed to reduce the risk-exposure.

VI. I-AM for eHealth Applications

To demonstrate the applicability of the I-AM cycle, we describe how the three stages can be applied in the context of mobile eHealth applications, using the four apps introduced earlier.

A. Inform Strategies

It is essential that users are informed of all aspects that are potential causes for privacy concerns, especially because eHealth apps are handling particularly sensitive data about their users. Specifically, users need to be made aware of the data collected by an app, where and how this data is stored, when and to who it is transmitted, and how it may be used.

f) Status Quo in eHealth Apps: Unfortunately, for the apps we analyzed, this information is mostly buried in lengthy privacy policies that are full of jargon. In addition, some policies promise to never share user data at the beginning of the policy, but then qualify this much further down in the policy, allowing them to share non-personally identifiable information (PII) with arbitrary third parties. This is especially problematic because it has been shown that re-identification of non-PII is possible [21]. Another problem concerns the way that requested permissions are displayed to the user. As previous work has pointed out, these are usually only shown after the user has made the decision to install an app. In addition, some apps request permissions that seem unnecessarily broad, or that are surprising because it is not obvious why this permission would be needed to provide the app's functionality.

g) Inform Options: There is a body of existing research on how to better inform users about privacy aspects. In the context of mobile applications, several papers have discussed ways to alter the app store design to improve when and how requested permissions are presented to users [5], [6].

To improve the usability of privacy policies, researchers have advocated for machine-readable privacy policies that could then be matched automatically against a user's privacy objectives [22]. However, work that addresses the issue of interface design for this approach is sparse and focuses on interfaces for PCs, not on mobile devices [23].

h) Discussion on Impact: Improvements of user information generally do not impact on an app's functionality, because this stage takes place before an app is installed by the user. However, the improvements suggested by previous work do influence users' privacy awareness and installation decisions [5], [6], which we see as evidence that such improvements are much needed.

B. Alerting Strategies

Once an app has been installed, the Alert stage notifies users of current and ongoing privacy risks. This is an essential stage to increase privacy awareness and is a prerequisite for mitigation.

i) Status Quo in eHealth Apps: The apps we analyzed do not alert users to currently ongoing risks that could cause privacy concern. In particular, users are not alerted to transfers of data, including transfers to third parties. The only alert with regard to the collection of data is provided by the Android operating system, which shows a symbol in the phone's status bar whenever an app requests detailed location information via GPS. This is clearly an unsatisfactory situation from a user's point of view, and it highlights that today's apps regard the publication of a privacy policy as sufficient fulfillment of their privacy-related duties.

j) Alert Options: The first consideration for the alert stage is what alerts should be about. In general, alerts combine information from the Inform stage with contextual information, and thus determine the severity of the resulting privacy risk. Contextual information includes ongoing data transfers as well as access to operating system permissions such as location or camera. The severity of privacy risks can be determined from user preferences, crowd-sourced ratings of surprise [6], or technical indicators such as unencrypted data transfers.

The second consideration is the technical realization. For users, the easiest option would be an integration of alerting into each app. If app developers can be motivated to include this feature at all, it would take a long time until all apps were modified. Therefore, we believe that alerting may better be realized as a separate layer located between apps and the operating system. This layer could record permission usage and analyze data transfers, and alert the user in a consistent manner for many different apps.

The third consideration is how to design user interfaces for alerting. A simple option would be to display a pop-up window for each alert. However, this is likely to happen frequently, and thus likely to annoy users and disrupt their app usage. Another option would be to display small icons in the device's status bar, similar to those used by Android to indicate GPS usage. The problem with these icons is that they are extremely non-obtrusive and thus may easily go unnoticed. One could imagine several alternative ways to display alerts which would be located between these two extremes in terms of obtrusiveness, for example colored window borders, colored backgrounds, tooltips, or even non-visual cues such as vibration.

We believe that the most promising approach to designing a user interface for alerting is a scalable interface, where the prominence of an alert depends on a user's privacy preferences or objectives. To achieve such a scalable interface, the severity

of an alert is determined based on the user's privacy preferences. For each level of severity, a different method will be used to display the alert. For example, if a user is not concerned about an app's usage of coarse location information, the alert would be displayed as a small icon. On the other hand, if a user is very concerned about apps accessing his contact list, the alert would be presented as a large pop-up window.

k) Discussion on Impact: Depending on how alerts are realized, they may interfere with how a user wishes to use an app. For example, too-frequent pop-ups may become annoying and may cause the user to disable the alerting feature entirely. Therefore, scalable interfaces need to be designed carefully. We believe a scalable interface should offer configuration options to users, and should integrate machine learning so that the interface can adapt itself to a user's past choices. In addition, care must be taken to present privacy-related information in an easy-to-understand, non-technical way. If these issues are taken into account, a well-designed, adaptive interface will not substantially interfere with app functionality after a short period of configuration and learning.

C. Mitigation Strategies

After users have been informed of potential privacy issues, and alerted to ongoing real privacy issues, mitigation strategies aim to give users options for dealing with privacy issues.

l) Status Quo in eHealth Apps: The four eHealth apps we analyzed in this paper do not offer many mitigation options. All apps protect the privacy of data in transit against network eavesdroppers by encrypting data transfers with SSL/TLS. Aside from this, the only mitigation option we found is one app's privacy policy which advises users that they may uninstall the app if they disagree with the data usage/storage/sharing policies.

m)Technical Mitigation Options: From a technical perspective, it is easy to imagine a range of mitigation options that do not involve uninstalling an app. In some cases, it may be possible for users to switch to an open source app that may have a limited set of features, but may also have fewer privacy issues. The features and privacy issues we highlighted for Runtastic and RunnerUp in Sections III and IV may serve as an example for this switch. In case switching apps is not a desirable mitigation option, it is possible to insert a privacy-enhancing layer located between operating system and app. This layer could provide users with a range of different mitigation options, for example it could offer to:

- store data locally by intercepting data transfers to the app provider's servers
- aggregate data, for example send a daily summary of calorie intake, but not the individual foods consumed
- filter data, for example remove sensitive heart rate data from activities sent to a fitness tracker
- transform data, for example alter the GPS coordinates in a recorded activity to conceal where the activity took place.

These technical mitigation options need to be integrated into the user interface in a way that they empower the user to make use of mitigation, but do not distract from the user's main task or annoy by too-frequent interactions. The easiest way to realize this is to integrate the mitigation interface into our proposed scalable interface for alerting. In this way, mitigation options will be presented differently depending on the severity of the current privacy risk. Similar to the alert interface, the mitigation interface should be able to learn from a user's past choices. In this way, a user's favorite mitigation options can be offered in a convenient manner, while less-frequently used choices can be harder to reach. In addition, it should be possible to configure automatic mitigation actions for specific situations.

Mitigation options do not necessarily have to be retrofitted with a separate privacy-enhancing layer, they can also be integrated directly into an app. For example, RunnerUp offers users the choice if and when to upload an activity to a third-party website. This combines a default to local storage with empowering the user to control selective data sharing.

n) Discussion on Impact of Mitigation: It is obvious that some mitigation options will have an impact on the functionality and usability of apps. Mitigation by uninstallation is clearly disruptive, because it removes access to the app's functionality. The other mitigation options are less disruptive, but may still reduce functionality, depending on the specific mitigation technique. This reduction needs to be balanced with the user's desire to access certain features. By presenting users with different mitigation options, I-AM empowers users to control this balance themselves.

VII. Conclusion

We have presented the Inform–Alert–Mitigate (I-AM) cycle, a structured tool to address users' privacy concerns when using mobile applications. We demonstrated the effectiveness of I-AM with a case study of four eHealth applications, including two fitness trackers, a dieting app, and a medical app for diabetes patients. We showed how a systematic consideration of alerting and mitigation options can improve user privacy while preserving functionality, and give control over sensitive data back to users.

In the broader context, we see our work embedded in Responsible Research and Innovation (RRI) [24], a framework that aims to ensure desirable and acceptable outcomes of research and innovation. Our paper, and the I-AM cycle we present, is one step towards making app design more responsible by providing developers with a tool for systematic assessment and improvement of privacy aspects.

In the future, we plan to evaluate the increase in privacy awareness and privacy protection that can be achieved by following the I-AM approach, both theoretically (applying our previous work on privacy measurement [25], [26]) and with user studies. In addition, we plan to investigate how I-AM can be applied in security scenarios, especially as it concerns mobile workers who need to balance their productivity with their organization's effective security, especially as their work is done in a variety of physical spaces each with different privacy and security constraints.

References

[1] G. Greenleaf, "Sheherezade and the 101 Data Privacy Laws: Origins, Significance and Global Trajectories," *Journal of Law, Information and Science*, vol. 23, no. 1, pp. 4–46, 2014.

[2] A. P. Felt, S. Egelman, and D. Wagner, "I've Got 99 Problems, but Vibration Ain't One: A Survey of Smartphone Users' Concerns," in *Proceedings of the Second ACM Workshop on Security and Privacy in Smartphones and Mobile Devices*, ser. SPSM '12. New York, NY, USA: ACM, 2012, pp. 33–44.

[3] G. Iachello and J. Hong, "End-user Privacy in Human-computer Interaction," *Found. Trends Hum.-Comput. Interact.*, vol. 1, no. 1, pp. 1–137, Jan. 2007.

[4] A. S. Patrick and S. Kenny, "From Privacy Legislation to Interface Design: Implementing Information Privacy in Human-Computer Interactions," in *Privacy Enhancing Technologies*, ser. Lecture Notes in Computer Science, R. Dingledine, Ed. Springer Berlin Heidelberg, 2003, no. 2760, pp. 107–124.

[5] P. G. Kelley, L. F. Cranor, and N. Sadeh, "Privacy As Part of the App Decision-making Process," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '13. Paris, France: ACM, 2013, pp. 3393–3402.

[6] J. Lin, S. Amini, J. I. Hong, N. Sadeh, J. Lindqvist, and J. Zhang, "Expectation and Purpose: Understanding Users' Mental Models of Mobile App Privacy Through Crowdsourcing," in *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, ser. UbiComp '12. Pittsburgh, PA, USA: ACM, 2012, pp. 501–510.

[7] F. Schaub, R. Balebako, A. L. Durity, and L. F. Cranor, "A Design Space for Effective Privacy Notices," in *Symposium On Usable Privacy and Security (SOUPS 2015)*. Ottawa, Canada: USENIX Association, Jul. 2015.

[8] D. Rosenberg, S. Foley, M. Lievonen, S. Kammas, and M. J. Crisp, "Interaction spaces in computer-mediated communication," *AI & SOCIETY*, vol. 19, no. 1, pp. 22–33, Nov. 2004.

[9] W. Enck, D. Octeau, P. McDaniel, and S. Chaudhuri, "A Study of Android Application Security." in *USENIX security symposium*, vol. 2, San Francisco, CA, USA, Aug. 2011.

[10] P. Hornyack, S. Han, J. Jung, S. Schechter, and D. Wetherall, "These Aren't the Droids You're Looking for: Retrofitting Android to Protect Data from Imperious Applications," in *Proceedings of the 18th ACM Conference on Computer and Communications Security*, ser. CCS '11. Chicago, IL, USA: ACM, Oct. 2011, pp. 639–652.

[11] A. P. Felt, E. Ha, S. Egelman, A. Haney, E. Chin, and D. Wagner, "Android Permissions: User Attention, Comprehension, and Behavior," in *Proceedings of the Eighth Symposium on Usable Privacy and Security*, ser. SOUPS '12. Pittsburgh, PA, USA: ACM, Jul. 2012, pp. 3:1–3:14.

[12] C. Bonnington, "Apple Says Grabbing Address Book Data Is an iOS Policy Violation," Feb. 2012. [Online]. Available: http://www.wired.com/2012/02/apple-responds-to-path/

[13] S. Egelman, J. Tsai, L. F. Cranor, and A. Acquisti, "Timing is Everything?: The Effects of Timing and Placement of Online Privacy Indicators," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '09. Boston, MA, USA: ACM, Apr. 2009, pp. 319–328.

[14] P. G. Kelley, S. Consolvo, L. F. Cranor, J. Jung, N. Sadeh, and D. Wetherall, "A Conundrum of Permissions: Installing Applications on an Android Smartphone," in *Financial Cryptography and Data Security*, ser. Lecture Notes in Computer Science, J. Blyth, S. Dietrich, and L. J. Camp, Eds. Springer Berlin Heidelberg, 2012, no. 7398, pp. 68–79.

[15] A. Barth, A. Datta, J. Mitchell, and H. Nissenbaum, "Privacy and contextual integrity: framework and applications," in *2006 IEEE Symposium on Security and Privacy*, Oakland, CA, USA, May 2006, pp. 15 pp.–198.

[16] L. Barkhuus, "The Mismeasurement of Privacy: Using Contextual Integrity to Reconsider Privacy in HCI," in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ser. CHI '12. NAustin, TX, USA: ACM, 2012, pp. 367–376.

[17] H. Xu, N. Wang, and J. Grossklags, "Privacy by ReDesign: Alleviating Privacy Concerns for Third-Party Apps," in *ICIS 2012 Proceedings*, Orlando, FL, USA, Dec. 2012.

[18] T. Fischer, A. Sadeghi, and M. Winandy, "A Pattern for Secure Graphical User Interface Systems," in *20th International Workshop on Database and Expert Systems Application, 2009. DEXA '09*, Linz, Austria, Aug. 2009, pp. 186–190.

[19] B. Falchuk and S. Loeb, "Privacy enhancements for mobile and social uses of consumer electronics," *IEEE Communications Magazine*, vol. 48, no. 6, pp. 102–108, Jun. 2010.

[20] G. Faden, "Multilevel Filesystems in Solaris Trusted Extensions," in *Proceedings of the 12th ACM Symposium on Access Control Models and Technologies*, ser. SACMAT '07. Nice-Sophia Antipolis, France: ACM, 2007, pp. 121–126.

[21] A. Narayanan and V. Shmatikov, "De-anonymizing Social Networks," in *2009 30th IEEE Symposium on Security and Privacy*, May 2009, pp. 173–187.

[22] A. I. Antón, E. Bertino, N. Li, and T. Yu, "A Roadmap for Comprehensive Online Privacy Policy Management," *Commun. ACM*, vol. 50, no. 7, pp. 109–116, Jul. 2007.

[23] J. Angulo, S. Fischer-Hübner, E. Wästlund, and T. Pulls, "Towards usable privacy policy display and management," *Information Management & Computer Security*, vol. 20, no. 1, pp. 4–17, Mar. 2012.

[24] B. C. Stahl, "Responsible research and innovation: The role of privacy in an emerging framework," *Science and Public Policy*, vol. 40, no. 6, pp. 708–716, Dec. 2013.

[25] I. Wagner and D. Eckhoff, "Privacy Assessment in Vehicular Networks Using Simulation," in *Proc. Winter Simulation Conf. (WSC '14)*, Savannah, GA, USA, December 2014.

[26] I. Wagner, "Genomic Privacy Metrics: A Systematic Comparison," in *2015 IEEE Security and Privacy Workshops (SPW)*, San Jose, CA, May 2015, pp. 50–59.