Privacy, Security, and Machine Learning for Mobile Health Applications

Lucila Ohno-Machado[1], Shuang Wang[1],
Xiaofeng Wang[2], Arya Iranmehr[1], Xiaoqian Jiang[1]

[1]Division of Biomedical Informatics
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093, USA

[2]School of Informatics and Computer Science
Indiana University at Bloomington
150 S. Woodlawn Avenue
Bloomington, IN 47405-7104

Contact: lohnomachado@ucsd.edu

Abstract

Mobile health (mHealth) applications have become very popular in the past few years due to improvements in hardware, telecommunications, software applications ("apps"), and affordability of devices and data transmission plans. mHealth apps can now use several sensors to assess an individual's health status and assist with health-related decision making. The increase in mHealth applications has not been paralleled by an increase in mechanisms to protect data integrity and individual privacy. In this article, we provide an overview of the main types of security and privacy threats, as well as constraints faced by machine learning algorithms that operate on mHealth devices.

mHealth devices and applications

A recent study conducted by the IMS institute [1] evaluated 43,689 apps that were listed under the category 'healthcare and fitness' at the Apple App store by June 2013. Of these, 16,275 apps were consumer- or patient-oriented. Many others target healthcare providers. Furthermore, 5,095 apps had the capability to capture data entered by users and 395 apps could communicate with healthcare providers or share data in social networks. There were 159 apps that were able to connect with external sensors, although fewer than 50 could measure vital signals. Table 1 summarizes a few top apps categorized by their functionalities.

Table 1. Example of various top apps and their features

App Name	Description	Record personal information	Access internal sensors	Access external sensors	Send data outside the smartphone
ZocDoc	Find & book doctor appointments	X	location		X
Healow	Communicate with doctor's office and accessing medical records	X	location		X
Calorie Counter	Calorie counting app with a large food database	X	location, camera		X
Strava	Workout tracking and activity analysis	X	location, microphone	Bluetooth sensors	X
Healthtap	Health related Q&A from more 62,000 U.S. doctors	X	location, camera, microphone	Bluetooth sensors	X
CVS pharmacy	Refill, transfer and prescription management	X	location, camera		X
Dosecast	Medication reminders and medication adherence tracking	X			
Glucose Buddy	Data storage utility for diabetes patients	X			
Baby Connect	Baby care tracking application	X	location		
OneTouch Reveal	Blood sugar data checking from OneTouch Verio Sync glucose meter	X		Bluetooth sensors	

In a similar way, we summarized the top 5 apps in the iOS category 'Health & Fitness' (as of 9/29/2014) in Table 2.

Table 2. Summary of top grossing apps for 'Health & Fitness' on Apple's App Store

App Name	Description	Record personal information	Access internal sensors	Access external sensors	Send data outside smartphone
Lose it!	Social network for weight loss program and calorie counter	X	location		X
7 Minute workout challenge	A research-backed workout program	X			X

Sleep cycle alarm clock	An intelligent alarm clock that wakes people up in the lightest sleeping phase	X		Bluetooth sensors	X
RunKeeper	App that tracks pace and measures workout distance	X	location	Bluetooth sensors	X
FitStar personal trainer	Fitness app that crafts a personalized workout plan	X			X

Our tables just sampled a small portion of the mobile health app market, which is expanding very quickly. Mobile health apps use various sensors and mechanisms to collect heterogeneous personal health information and interact with users. The policies for re-use of data are not always evident and it is unclear whether users fully understand risks and benefits. Among many critical aspects, we found that security, privacy, and learning issues are dealt with great diversity among the apps.

Security: protecting data and application integrity

Wireless connectivity (e.g., Bluetooth, WiFi, NFC, etc.) provides great flexibility for mobile health (mHealth) devices or apps to communicate with other terminals (e.g., healthcare providers, service clouds, or other mHealth sensors) or with a mobile gateway, which can interact with remote users and platforms like Google Health. Figure 1 depicts a Body Area Network (BAN), which consists of short-range wireless sensors including both wearable and implanted devices. Mobile sensors within the same BAN can wirelessly communicate with a mobile gateway (e.g., a smartphone) or other sensors. Moreover, the mobile gateway can wirelessly connect with remote servers for the purposes of storing health data in a cloud, monitoring a particular health condition or vital signs

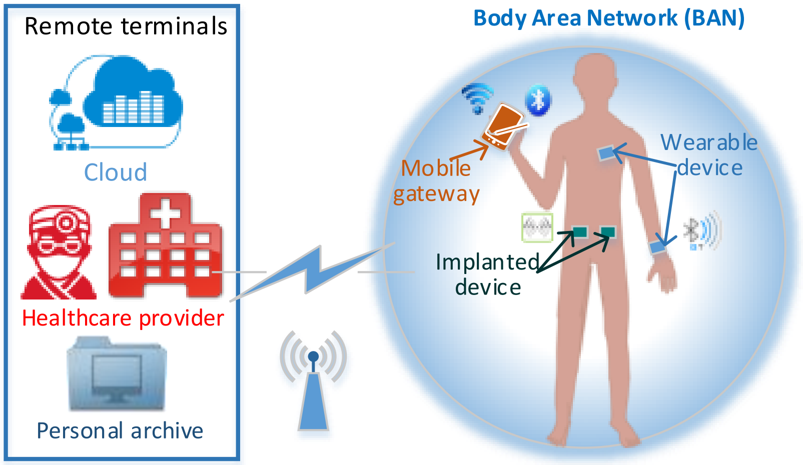

Figure 1: BAN is a short-range (e.g., a few feet) wireless sensor network including both wearable or implanted devices, which can communicate with each other or with a mobile gateway to exchange physiological and environmental data, such as *blood pressure, glucose level, temperature, heart reads, moving speed, body posture*, etc. A BAN can also interact with remote users and platforms such as *cloud, healthcare provider, or home health information archive*.

in an emergency, and archiving health data in a personal computer. For example, Moulton *et. al.* [2] show the feasibility of wirelessly uploading electronic health records directly from sensor systems. One benefit of importing data into the EHR directly from mobile sensors could stem from the minimization of certain manual data entry processes, as these processes are usually costly and may not always be reliable due to potential human errors. However, in addition to potential benefits, the wireless communication from the mHealth devices can also impose challenges and increase vulnerabilities related to data security and privacy. Table 3 summarizes the threats, risks and corresponding defense mechanisms in mHealth applications.

Resource depletion (RD) attack: A RD attack [3] attempts to exhaust resources of an mHealth device, such as, battery, bandwidth, storage and so on. This is similar to denial of service (DoS) attacks [4] on the internet, which cause a service or network resource to become unavailable to the intended users. Due to portability requirements, mHealth devices usually have somewhat limited battery life, storage, computational power and communication bandwidth. Importantly, an mHealth device should only be able to communicate with an authenticated reader, e.g., the patient's healthcare provider. The authentication step between a reader and a mHealth device requires both communication and computation operations, which consume a considerable amount of battery and communication bandwidth. In addition, for security purposes, each authentication activity needs to be logged in the mHealth device. Therefore, if a malicious reader constantly requests authentication from an mHealth device, this can result in extremely fast battery consumption as well as allocation of a considerable amount of logging storage. For example, an insulin pump depends on the reads from wireless glucose sensor to accurately deliver the correct insulin dose to a diabetes patient. An RD attack can reduce the effective standby time of the glucose sensor from several months to a few hours, which may have serious consequences.

To mitigate an RD attack, an mHealth device may adopt a limited authentication protocol instead of an on-demand based communication strategy. For example, an authentication request should be only allowed within certain locations (e.g., patient's home or office), or during fixed time windows (e.g., every 30 minutes), or when certain patient conditions are observed (e.g., extreme glucose readings). The above strategies rely on the development of advanced access control mechanisms [5–7].

Replay attack: A replay attack attempts to spoof the sensor readings to induce the users or other mHealth devices to make wrong decisions [8]. Before establishing wireless communication links among mHealth devices, it is usually necessary to pair devices with a unique passcode (e.g., the serial number of a device) to enhance security. Data encoding and transmission formats must satisfy security concerns. That is, without knowing the passcode and data encoding/transmission format, it should be difficult to spoof the sensor data directly. However, even without knowing the meaning of the underlying data, an attack can still record the communication among mHealth devices over the air, as all data are transmitted wirelessly in a beacon fashion (i.e., communications are repeated in a fixed interval). A valid data transmission could be maliciously repeated by a hacker. For example, an attacker can first record the communication from the sensor that indicates a high glucose level based on some auxiliary knowledge about the victim. Then, the attacker could later retransmit the high glucose information pretending it is a "valid" message, which would cause the receiver (e.g., insulin pump) to deliver an incorrect insulin dose and put the patient at risk.

One way to avoid replay attacks is to introduce timestamps [9] within the message, where one mHealth device only accepts the message from the other device if the timestamp in the received message is within a reasonable time tolerance range. In this case, a replay attack would not be able to provide a valid timestamp by simply reusing the previously sniffed transmission. However, timestamping requires synchronization with mHealth devices, which may impose additional communication burden and reduces battery life.

External Device Mis-Bonding (DMB) attack: A DMB attack [10] targets the mobile gateway running the Google Android Operating System (OS). Android devices have gained 85% market share in Q2 2014 [11]. Android allows external mHealth devices to connect with the smartphone by sharing the same channel using Bluetooth, WiFi, NFC, etc. Many apps in Android devices can request permission to use this channel for different purposes. Consequently, any app that has been granted permission on the communication channel for a given purpose could act as an insider to steal sensitive information from another app running on the same channel [10]. This is because Android only ensures authorized connections among sensors and the smartphone instead of among the authorized apps residing in the smartphone. In particular, there are two security issues associated with the DMB attack, including the information leakage risk and the information injection risk. For the first type, a malicious app on an authorized phone can steal sensitive patient data that are intended to be transmitted to an authorized app. In terms of information injection risk, an authorized app can feed false medical information into the original authorized app by intercepting the connection between the authorized external device and the authorized smartphone. This injection risk is extremely dangerous for patients who are heavily relying on health monitoring apps (e.g., blood-sugar concentration, irregular heart rhythm, etc.). A recent study [10] shows that four popular mHealth devices including iThermometer [12] (real time body temperature monitor), as well as FDA approved Class II medical devices such as Bodymedia FIT Link Armband [13] (daily activity monitor), Nonin Onyx II Pulse Oximeter [14], and Entra Health System Blood Glucose Meter [15] lack the secure bonding control between the devices and their official apps. Moreover, the study also analyzed 68 official Android apps with external device connection capabilities and confirmed that none of these apps have been equipped with app-to-device authentication mechanisms.

To tackle this type of threat, an OS-level safeguard mechanism is needed, which require an Android system update from Google, App improvement by developers, as well as mHealth device manufacture enhancements. A protection app called Dabinder [10] was developed to block malicious attempts to bind an external mHealth device. However, such an app needs to continuously run in background, which may affect the performance of the smartphone.

Table 3: Summary of threats, risks and defense mechanisms in mHealth applications

Threat	System/Device	Risk	Defense
RD attack	Pacemaker, implantable cardioverter defibrillator, insulin pump, glucose sensor, intrathecal drug delivery, heart rate monitor, etc.	shorter effective lifetime	access control, shorter range communication
Replay attack		malfunction, irregular actions	time-stamping
DMB attack	Smartphones and their external mHealth devices	information leak, information injection	App-to-device level authentication

Privacy: protecting health information

The Health Insurance Portability and Accountability Act of 1996 (HIPAA) [16] can be extended to mHealth. Pursuant to HIPAA, the U.S. Department of Health and Human Services (HHS) issued a regulation governing the privacy of health information that is maintained in electronic form, known as the “Privacy Rule” [17]. The rule requires covered entities (or a business associate/service provider of such covered entities) that transmit health information to ensure the confidentiality of the information. This information should not be disclosed except for the purposes of treatment, payment for treatment, and health care operations without authorization from the patient or study participant. HIPAA allows healthcare providers to communicate electronically with patients but it requires “appropriate administrative, physical and technical safeguards to ensure the confidentiality, integrity, and security of personal health information [18].” We summarize the safeguards related to mHealth in Figure 2.

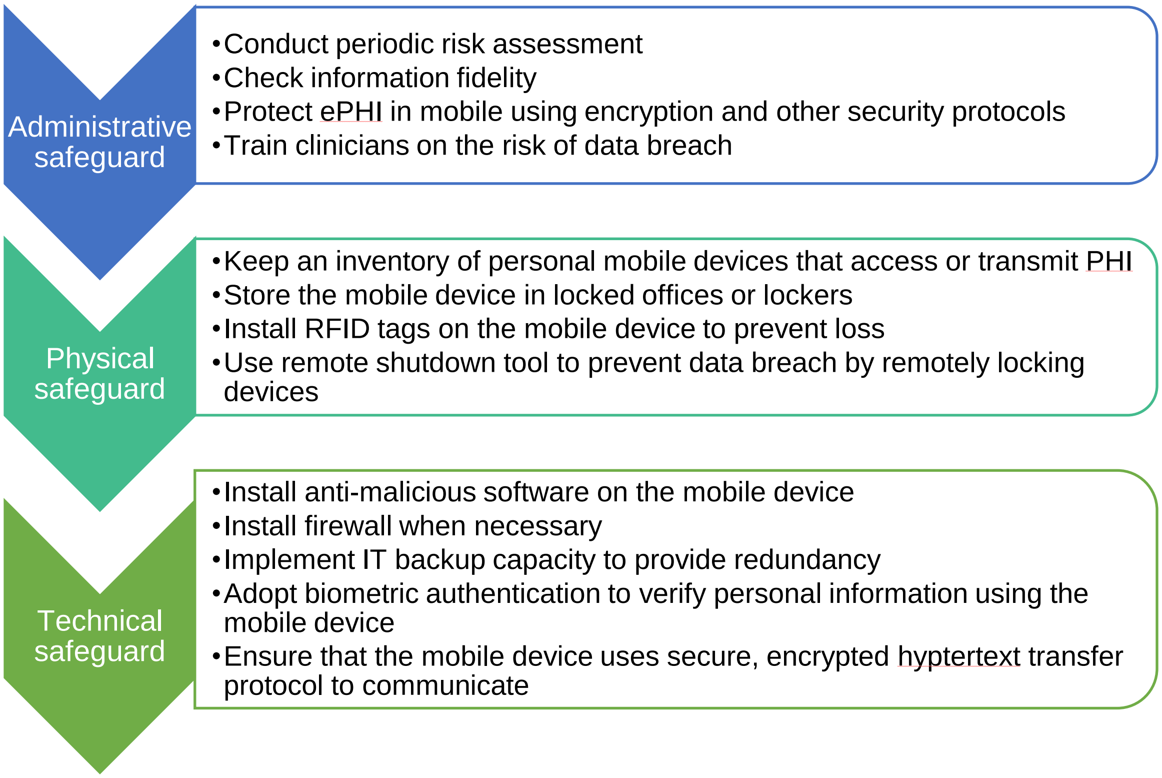

- Conduct periodic risk assessment
- Check information fidelity
- Protect ePHI in mobile using encryption and other security protocols
- Train clinicians on the risk of data breach

- Keep an inventory of personal mobile devices that access or transmit PHI
- Store the mobile device in locked offices or lockers
- Install RFID tags on the mobile device to prevent loss
- Use remote shutdown tool to prevent data breach by remotely locking devices

- Install anti-malicious software on the mobile device
- Install firewall when necessary
- Implement IT backup capacity to provide redundancy
- Adopt biometric authentication to verify personal information using the mobile device
- Ensure that the mobile device uses secure, encrypted hyptertext transfer protocol to communicate

Figure 2: Three types of safeguards pertaining to mobile devices [19].

However, HIPAA’s privacy rule only applies to mHealth applications that involve both a covered entity (healthcare providers like hospitals, clinics, and doctors, but not personal health service providers, such as Google Health) and protected health information (PHI). For example, most existing health and fitness apps, which do not transmit users’ PHI to a healthcare provider, are not subject to the HIPAA privacy rule. There are other applicable laws like the Children’s Online Privacy Protection Rule (COPPA) [20], the Federal Trade Commission Act (FTC Act) [21], and state laws (e.g., for breach notification). Each governing rule has its own scope. Table 4 lists the obligated entities and data covered by different laws.

Table 4: Summary of obligated entities and data covered by different laws.

	Obligated entities				**Covered data**			
	Covered entities	Operators who manage personal data	Legal entity owns information of state residents	All commercial entities	Protected health information	Children's personal information	Consumer data related to identity	"Personal information" at large
HIPAA	X				X			
COPPA		X				X		
FTC Act	X	X	X	X			X	
state laws			X					X

Existing laws do not fully cover patient health information created in mHealth and there are worries that device vendors, mobile operators, and health app owners may have access to health care data that patients have little knowledge of or control over [22]. A recent article surveyed availability and quality of mHealth app privacy policies, showing that current policies do not make information privacy practices transparent to users, require college-level literacy, and are often not focused on the app itself [23].

We summarize some of the major technical issues that deserve attention.

- Anonymity for personal health data: personal health data can still be vulnerable after removal of explicit identifiers such as name and social security number. Sweeney shows that 87% of the U.S. population can be uniquely re-identified using a triplet of gender, date of birth, and zip code [24]. A simple attack can be achieved by linking these attributes to public databases such as the voter registries.

- Location service: many mHealth apps use location information to log workout and provide context aware services. These services introduce privacy risks because the history of location is very sensitive and can be abused by adversaries if not protected appropriately. Location privacy is an active area of research [25,26]. Most recently, Lee et al. proposed a novel way to hide location semantic information by cloaking it with geographically close but semantically heterogeneous locations [27].

- Secure data transmission: attackers might be able to acquire sensitive information over the wireless network by monitoring unencrypted packets. Users of the mHealth app need to be aware of whether the device communicates with the server using secure protocols with encryption and whether it stores raw data in a SIM-card, SD card, or other media.

- Device presence: wireless communication between mobile and sensor (e.g., device pairing using Bluetooth or WiFi protocols) might allow others to discover what type of medical sensor an individual is carrying. This can be disadvantageous in some situations. Singelee and Preneel developed a secure distance bounding protocol to cloak device presence [23], but it has not yet been adopted in real apps.

For more detailed privacy issues in mobile technology for personal healthcare, interested readers are referred to the survey of Avancha and Baxi [22], which covers identification threats, authentication, consent management, access control, and auditing.

Machine learning methods: types and constraints

Mobile and wearable devices such as smartphones, smartwatches, glasses, etc. are now available with a wide range of sensors (e.g. vision, audio, light, temperature, magnetism, pressure, proximity, direction, position, acceleration, etc.) at a reasonable price. Ubiquity of mobile devices has enabled machine learning algorithms to be applied to process sensors' streaming multi-view data in various tasks such as:

I **Intervention.** Although mHealth has shown some examples of how it can be effective in health and behavioral interventions [28], most applications are yet to implement efficient, in-time, adaptive, context-aware and interactive interventions [29]. Such complex interventions, which require inference, cannot be done in large scale without machine learning tools.

II **Recommendation.** Recommender systems have been widely studied in the machine learning community [30] and can be utilized in mHealth in a such way that users receive health recommendations according to their preferences and their similarity to other users.

III **Evidence Generation.** mHealth could leverage statistical machine learning to evaluate efficacy of interventions [31] and analyze time series data to perform:

- **Monitoring.** Continuous monitoring at individual or population levels could be done for determining trends and detecting anomalies or outliers.
- **Prediction.** Given the computational power of recent mobile devices, real-time prognosis, outcome prediction, survival analysis are now possible on mobile devices using current machine learning models and user feedback. Retrospective analyses may also be useful to identify risks and protective factors for a user or population.

However, key challenges to conduct the above tasks exist.

1. **Supervised learning.** The learning phase of supervised learning algorithms usually involves processing a labeled dataset, i.e., a set with a 'gold standard'. In practice, for the training data, class labels are assigned to samples by humans (e.g. expert or crowdsourcing) in a costly and time-consuming operation that is not feasible for the multi-modal and streaming data in mHealth. Although the (occasional and noisy) user feedback provides labeled training samples, collecting a large enough labeled dataset still remains a challenge.

2. **Context-aware mHealth.** Performing intervention, recommendation and prediction with respect to context, i.e., time, location, social environment, psychophysiological status or mood is a difficult task not only because of sampling and algorithmic complexity but also because of the difficulties in understanding context. The current tool for understanding context in mHealth has been activity recognition [32,33]. For example, using activity recognition, an Ambulation app [34] enables doctors to monitor mobility and recovery of mobility-affecting chronic diseases (e.g. multiple sclerosis, Parkinson's).
3. **Multi-modal data.** The problem of learning from sensory data is highly multi-modal. Each modality (i.e. sensor, input), has its own statistical properties, e.g. noise levels, and contains different kinds of information. The key challenge of multi-modal machine learning is to combine multiple input channels (i.e., multi-sensor information fusion [35]) in order to provide a less noisy, more complete, consistent and reliable perception of the environment for *a specific task*. Such a heterogeneous and pervasive sensing for mHealth is likely to be addressed by deep learning methods in the near future [36].
4. **Personalized mHealth applications.** Individualized information could be delivered using different machine learning techniques.
    - **Online learning**. Learning a single static machine learning model for all users not only does not provide personalized prediction, but may also result in inferior accuracy when compared to dynamic and adaptive models [37]. This problem is a special case of online machine learning [38], in which, at each iteration, the learning model (1) receives an unlabeled sample, (2) makes a prediction, (3) receives the true label (probably by user feedback), and (4) updates the model.
    - **Incorporating prior knowledge.** Bayesian machine learning addresses this aspect. Theoretically, any *plausible* prior knowledge about the problem improves the prediction performance of the learning algorithm [39]. Intuitively, in the context of mHealth, the use of prior knowledge such as age, risk-factor, or even genomic data may help personalize the learning model and improve predictions for each user.
5. **Hardware resources.** Perhaps the so-called memory wall and power wall [40] that have been important for Moore's law [1] are still applicable in mHealth because the following power and memory intensive operations make them different from other applications:
    - Sampling data at high rates from different sensors
    - Complexity of the machine learning model
    - Privacy preserving operations
    - Encryption and transmission to servers
    - Classification of new samples in real-time
    - Online learning that requires constantly updating the learning model

[1] The performance of processors doubles every 18 months

Summary and Future Directions

In summary, the number and nature of mHealth applications has grown so rapidly in the past few years that a corresponding security and privacy framework has not yet been fully developed to ensure the integrity of the data against various types of attacks and the protection of individual privacy. Perhaps because there have been no serious publicized privacy or security breaches, hardware manufacturers, app developers, users, and health organizations have not yet demanded the development of specific policies. This development could not only provide reassurances for users, but could also encourage the modification of current policies that may have become obsolete or unnecessary, as users can increasingly customize their privacy and security settings. We are living in an exciting era in which users are technically able to exert more control over the security of mHealth devices and privacy of mHealth apps, and over the use of their protected health information. As we develop the policy framework for protecting privacy and ensuring integrity of mHealth data, we will learn important lessons that may carry on to Health data in general.

Acknowledgments

This research is supported in part by NIH grants U54HL108460, R00LM011392, R21LM012060, and K99HG008175) and PCORI contract CDRN-1306-04819.

Reference

1 Aitken M, Gauntlett C. Patient apps for improved healthcare from novelty to mainstream. *Parsippany, NJ IMS Inst Healthc Infomatics* 2013.

2 Moulton B, Chaczko Z, Karatovic M. Updating electronic health records with information from sensor systems: considerations relating to standards and architecture arising from the development of a prototype system. 2009.

3 Hei X, Du X. *Security for Wireless Implantable Medical Devices*. Springer 2013.

4 Mirkovic J, Dietrich S, Dittrich D, *et al. Internet Denial of Service: Attack and Defense Mechanisms (Radia Perlman Computer Networking and Security)*. Prentice Hall PTR 2004.

5 Denning T, Fu K, Kohno T. Absence Makes the Heart Grow Fonder: New Directions for Implantable Medical Device Security. In: *HotSec*. 2008.

6 Rasmussen KB, Castelluccia C, Heydt-Benjamin TS, *et al.* Proximity-based access control for implantable medical devices. In: *Proceedings of the 16th ACM conference on Computer and communications security*. 2009. 410–9.

7 Allen M. Pacemakers and implantable cardioverter defibrillators. *Anaesthesia* 2006;**61**:883–90.

8 Radcliffe J. Hacking Medical Devices for Fun and Insulin : Breaking the Human SCADA System.

9 Tang H, Liu X, Jiang L. A Robust and Efficient Timestamp-based Remote User Authentication Scheme with Smart Card Lost Attack Resistance. *IJ Netw Secur* 2013;**15**:446–54.

10 Naveed M, Zhou X, Demetriou S, *et al.* Inside Job : Understanding and Mitigating the Threat of External Device Mis-Bonding on Android. 2014;:23–6.

11 Smartphone OS Market Share, Q2 2014.

12 iThermometer. http://www.ithermometer.info/

13 BodyMedia FIT LINK Armband. http://www.bodymedia.com/Support-Help/BodyMedia-FIT-BW

14 Nonin Onyx II Pulse Oximeter. http://www.nonin.com/onyx9560

15 Entra Health System Blood Glucose Meter. http://www.entrahealthsystems.com/wireless.html

16 Health Insurance Portability and Accountability Act (HIPAA). http://www.hhs.gov/ocr/hipaa

17 Patient Privacy in a Mobile World: A framework addresses privacy law issues in mobile health. 2013.http://mhealthalliance.org/images/content/trustlaw_connect_report.pdf

18 Human Services. Nationwide Privacy and Security Framework For Electronic Exchange of Individually Identifiable Health Information. 2008;:1–12.http://www.healthit.gov/sites/default/files/nationwide-ps-framework-5.pdf

19 Barrett C. Healthcare Providers May Violate HIPAA1 by Using Mobile Devices to Communicate with Patients By. *2011* 2014;**8**:1–5.

20 COPPA - Children’s Online Privacy Protection. http://www.coppa.org/ (accessed 28 Sep2014).

21 Federal Trade Commission Act | Federal Trade Commission. http://www.ftc.gov/enforcement/statutes/federal-trade-commission-act (accessed 28 Sep2014).

22 Avancha S, Baxi A, Kotz D. Privacy in mobile technology for personal healthcare. *ACM Comput Surv* 2012;**45**:1–54. doi:10.1145/2379776.2379779

23 Singelee D, Preneel B. Location verification using secure distance bounding protocols. In: *IEEE International Conference on Mobile Adhoc and Sensor Systems Conference, 2005.* IEEE 834–40. doi:10.1109/MAHSS.2005.1542879

24 Sweeney L. Uniqueness of simple demographics in the US population. *LIDAP-WP4 Carnegie Mellon Univ Lab Int Data Privacy, Pittsburgh, PA* 2000.

25 Krumm J. A survey of computational location privacy. *Pers Ubiquitous Comput* 2008;**13**:391–9. doi:10.1007/s00779-008-0212-5

26 Minch RP. Privacy issues in location-aware mobile devices. In: *37th Annual Hawaii International Conference on System Sciences, 2004. Proceedings of the*. IEEE 2004. 1–10. doi:10.1109/HICSS.2004.1265320

27 Lee B, Oh J, Yu H, *et al.* Protecting location privacy using location semantics. In: *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '11*. New York, New York, USA: : ACM Press 2011. 1289. doi:10.1145/2020408.2020602

28 Krishna S, Boren SA, Balas EA. Healthcare via cell phones: a systematic review. *Telemed J E Health* 2009;**15**:231–40. doi:10.1089/tmj.2008.0099

29 Riley WT, Rivera DE, Atienza AA, *et al.* Health behavior models in the age of mobile interventions: are our theories up to the task? *Transl Behav Med* 2011;**1**:53–71. doi:10.1007/s13142-011-0021-7

30 Recommender Systems Handbook. http://www.springer.com/computer/ai/book/978-0-387-85819-7 (accessed 29 Sep2014).

31 Manda SK. International Journal of Advanced Research in Computer Science and Software Engineering Privacy Preserving Support for Mobile Health Care using Message Digest. 2013;**3**:97–102.

32 Burns MN, Begale M, Duffecy J, *et al.* Harnessing context sensing to develop a mobile intervention for depression. *J Med Internet Res* 2011;**13**:e55. doi:10.2196/jmir.1838

33 Donker T, Petrie K, Proudfoot J, *et al.* Smartphones for smarter delivery of mental health programs: a systematic review. *J Med Internet Res* 2013;**15**:e247. doi:10.2196/jmir.2791

34 Ryder J, Longstaff B, Reddy S, *et al.* Ambulation: A Tool for Monitoring Mobility Patterns over Time Using Mobile Phones. In: *2009 International Conference on Computational Science and Engineering*. IEEE 2009. 927–31. doi:10.1109/CSE.2009.312

35 Khaleghi B, Khamis A, Karray FO, *et al.* Multisensor data fusion: A review of the state-of-the-art. *Inf Fusion* 2013;**14**:28–44. doi:10.1016/j.inffus.2011.08.001

36 Srivastava N, Salakhutdinov R. Multimodal Learning with Deep Boltzmann Machines. In: *Advances in Neural Information Processing Systems*. 2012. 2222–30.

37 Longstaff B, Reddy S, Estrin D. Improving activity classification for health applications on mobile devices using active and semi-supervised learning. In: *Proceedings of the 4th International ICST Conference on Pervasive Computing Technologies for Healthcare*. IEEE 2010. 1–7. doi:10.4108/ICST.PERVASIVEHEALTH2010.8851

38 Shalev-Shwartz S. Online Learning and Online Convex Optimization. *Found Trends® Mach Learn* 2011;**4**:107–94. doi:10.1561/2200000018

39 Bishop CM, others. *Pattern recognition and machine learning*. springer New York 2006.

40 Asanovic K, Wawrzynek J, Wessel D, *et al.* A view of the parallel computing landscape. *Commun ACM* 2009;**52**:56. doi:10.1145/1562764.1562783