

Unintended Consequences of Wearable Sensor Use in Healthcare

Contribution of the IMIA Wearable Sensors in Healthcare WG

M. Schukat¹, D. McCaldin², K. Wang³, G. Schreier⁴, N. H. Lovell³, M. Marschollek⁵, S. J. Redmond³

¹ Department of Information Technology, NUI Galway, Ireland

² Graduate College of Engineering and Informatics, NUI Galway, Ireland

³ Graduate School of Biomedical Engineering, UNSW Australia, Sydney, Australia

⁴ Digital Safety & Security Department, AIT Austrian Institute of Technology GmbH, Graz, Austria

⁵ Peter L. Reichertz Institute for Medical Informatics, University of Braunschweig-Institute of Technology and Hanover Medical School, Hanover, Germany

Summary

Objectives: As wearable sensors take the consumer market by storm, and medical device manufacturers move to make their devices wireless and appropriate for ambulatory use, this revolution brings with it some unintended consequences, which we aim to discuss in this paper.

Methods: We discuss some important unintended consequences, both beneficial and unwanted, which relate to: modifications of behavior; creation and use of big data sets; new security vulnerabilities; and unforeseen challenges faced by regulatory authorities, struggling to keep pace with recent innovations. Where possible, we proposed potential solutions to unwanted consequences.

Results: Intelligent and inclusive design processes may mitigate unintended modifications in behavior. For big data, legislating access to and use of these data will be a legal and political challenge in the years ahead, as we trade the health benefits of wearable sensors against the risk to our privacy. The wireless and personal nature of wearable sensors also exposes them to a number of unique security vulnerabilities. Regulation plays an important

role in managing these security risks, but also has the dual responsibility of ensuring that wearable devices are fit for purpose. However, the burden of validating the function and security of medical devices is becoming infeasible for regulators, given the many software apps and wearable sensors entering the market each year, which are only a subset of an even larger 'internet of things'.

Conclusion: Wearable sensors may serve to improve wellbeing, but we must be vigilant against the occurrence of unintended consequences. With collaboration between device manufacturers, regulators, and end-users, we balance the risk of unintended consequences occurring against the incredible benefit that wearable sensors promise to bring to the world.

Keywords

Wearable; monitoring, physiologic; software; privacy

Yearb Med Inform 2016;73-86

<http://dx.doi.org/10.15265/IY-2016-025>

Published online November 10, 2016

as Runkeeper (<http://www.runkeeper.com/>). Other common wearable consumer devices, often operating as wireless peripherals of the wrist sensor, include shoe-worn sensors to accurately measure step count, and electrocardiogram (ECG) chest-bands to measure heart rate.

These wearable devices are sold to a consumer market, enabling users to track their activity levels and fitness training performance. These devices are intended to increase interest and performance in exercise and fitness, so there is an expected preventative healthcare benefit from their use; note, they are not marketed as medical/healthcare devices, but as fitness and wellbeing devices. While these fitness devices are intended to modify behavior for the better, there is potential for unintended perversions of behavior, depending on the user's compliance to the feedback and advice provided by the device.

Introduction

Wearable Fitness Sensors

The wearable sensor market has recently exploded. The total number of devices shipped worldwide (18.1 million) in the second quarter of 2015 was up 223% on the same quarter in 2014 [1]. Fitness and wellbeing device manufacturer Fitbit was the market leader with 24.3% of all shipments, followed by Apple, who released their long-anticipated smartwatch, taking 19.9%.

These wearable devices typically use various combinations of sensors such as accelerometers and gyroscopes to estimate characteristics of body movement (activity type, step count, cadences, etc.), magnetometers to estimate heading, barometers to estimate altitude, and GPS modules to track global position. Most commonly, the wearable device is a wrist-worn device, similar in form to a wristwatch; although smartphones also contain these same sensors and can serve as a similar wearable sensor through the use of software applications (apps), such

Wearable Medical Sensors

Regarding medical devices, although the Holter wearable ECG monitor has been in clinical use since the 1960s, other medical devices have only more recently made the leap to become fully wireless and wearable. Examples include wearable ambulatory oximeters, such as the Nonin WristOx2 (<http://www.nonin.com/WristOx3100>), or Sotera's ViSi Mobile device which measures ECG, pulse oximetry, and continuous non-invasive blood pressure, respiration rate, and skin temperature (<http://www.soterawireless.com/>).

Interestingly, there is an apparent convergence occurring between wearable fitness and wearable medical devices. The Apple Watch and the Samsung Gear are two examples of smartwatches which incorporate a photoplethysmogram (PPG) sensor to measure heart rate. Such technology is typically capable of measuring arterial blood oxygen saturation; although this function has not been made available by Apple yet, one would expect it will be a future feature. It is clear that these fitness devices are starting to encroach on the domain of wearable medical devices. This convergence of fitness and medical devices could see physiological data acquired using consumer electronics and interpreted by smartphone apps to provide medical advice. This trend will pose some significant challenges for regulators, such as the US Food and Drug Administration (FDA), whose task will be to ensure that these systems¹ are secure enough to protect the privacy of users and safe enough so as to not endanger the users' health.

The Internet of Things

While the scope of this paper is limited to unintended consequences that relate to the use of wearable sensors, it should be clear to the reader that many of these consequences are also relevant to the superset of devices that comprise the Internet of Things (IoT).

The IoT is defined as an extension to the current internet that enables connections and communication among "smart" objects. The International Telecommunications Union postulated this new type of ICT ecosystem in their 2005 report "The Internet of Things" [2], characterized by smart objects that interact with the physical world and communicate with each other around smart applications and services. Wearable sensors represent a subset of such smart objects, being able to identify, locate, sense, and connect, enabling communication between people and things.

¹ A system here comprising not only of the wearable device and its internal sensors, but of the communications network used to transmit the data, the Cloud storage where the data is housed, the software algorithms which interpret the data, and the interface which provides feedback and advice to the user.

The IoT has been identified as a technological solution to some personal connected healthcare challenges [3], and IoT-based architectures to facilitate healthcare-related applications have been proposed [4]; e.g., for the domain of Assisted and Healthy Aging in general [5], and for the management of particular chronic diseases, like COPD, in particular [6].

While IoT-based systems will enhance our ability to realize intended benefits for healthcare, it will most likely also amplify many of the unintended consequences discussed herein.

Overview of this Paper

In this paper we will discuss some of the unintended consequences, both beneficial and unwanted, resulting from the widespread adoption of wearable sensors by the general population. The paper is arranged around four broad topics where unintended consequences of the wearable sensor revolution could or already have become apparent:

1. Unintended modification of behavior;
2. Unintended creation of big data sets, and repercussion of their use and misuse;
3. Unintended privacy and security issues which are specific to wearable devices;
4. Unanticipated challenges facing regulatory bodies, which must somehow regulate both the safety and security of wearable devices and associated apps which interpret the acquired data.

Where possible, when an unwanted unintended consequence is highlighted, potential solutions will be proposed.

Unintended Modification of Behavior

A class of unintended consequences involves a scenario where sensors are used to provide patients or still healthy people (called 'subjects' in the following) with feedback on certain physiological parameters and/or aspects of their daily lives, to motivate them to change their behavior.

In many cases, wearable sensor systems are deployed to collect data which are supposed to support a lifestyle modification process. Examples are programs where chronically ill patients are equipped with sensors to track physical activity and thereby gauge their level of activity during daily life. A common concept is to use such sensors to monitor the impact of lifestyle change programs in a "before and after" paradigm.

Such systems typically either target the outcome parameter of such a program directly, e.g., the level of physical activity or the blood pressure, or, quite often, primarily address the adherence of the subject to a lifestyle modification or disease management program. The intervention being monitored by such methods may be the intake of medication, the frequency of gym visits, the behavior with respect to nutritional aspects, etc.

For some adherence-related programs, dedicated sensors and a body area network may be used; e.g., a patch on the belly to record the digestion of a pill [7]. In many cases, however, this type of monitoring may be facilitated just by tracking the location or classifying the level and type of activity using a smartphone carried by the subject.

Crucial elements of all those approaches are that the subjects receive feedback to trigger or enforce the learning mechanism, in the hope that this will lead to behavioral changes deemed necessary so as to improve the health or wellbeing of the subject. In prevention programs these measures aim to reduce the risk of future adverse events occurring. Since human behavior, however, is complex and occasionally irrational, people may react to this kind of feedback in a way not intended by the initiators of such programs.

One or more of the unintended consequences from the following non-exhaustive list may occur:

1. Subjects change their behavior for the worse

They may decrease their level of physical activity after getting the notion that they already are more active than they thought. Another aspect of this kind of behavioral modification is "trained helplessness". After some time being supported by technology, if this stimulus

is removed, subjects may no longer be able to sustain their improvements and revert to a level of activity/adherence below their levels on initiation into the program. For example, patients may lose their ability to cope with their overall medication plan once they have been “coached too intensively” by a supporting device (e.g., a wearable smart pill dispenser) and service that reminds them to take a particular medication in time.

2. Subjects may become more anxious about their health

They may develop a type of hypochondria or anxiety induced by this type of monitoring after being confronted with a health issue so intensively.

3. Subjects may become addicted to the device

Screen addiction to smartphones and computers is widespread [8, 9]. While wearable devices specifically for personal health monitoring are not as ubiquitous as generic personal mobile devices, it is logical that this addiction can be echoed in the wearable healthcare technology sphere, and users may become obsessed with self-monitoring beyond what can be considered a healthy level of attention to oneself. Individual reports of addictive behaviors regarding wearable fitness devices used by healthy individuals [10] provide insight into the unintended power that these devices yield and how they can shape the way users manage their daily life – it is *the tail wagging the dog*.

4. Subjects may adhere well to a program where adherence was expected to be poor

Somewhat similar to the scenario of addiction above is when subjects adhere perfectly to a program. An example of an unintended consequence in this case would be when a medication adherence management system (AMS) is used to ensure medications are taken according to schedule [11]. In a polypharmacy situation, however, patients are actually required to take a larger number of different drugs than they could reasonably manage. Patients often perform some self-optimization of the schedule based on their day-to-day experience in terms

of avoiding side effects. This may include changing the schedule or avoiding the intake of some of the drugs altogether which they feel impact poorly on their wellbeing. Introducing an AMS which enforces the prescribed medication plan may well cause serious side effects and subsequently decrease the level of adherence to a lower level than before. Alternatively, this could also lead to the patient consulting a doctor, or even lead to an unscheduled hospitalization, finally ending up with a more appropriate therapy regimen. In the end, this would be an unintended but positive consequence of such a system².

5. Subjects may fortuitously self-diagnose a problem

A similar unintended consequence may arise from people using wearable sensors for tracking various aspects of their lives in the context of the Quantified Self [12] which may initially not be health related. Analysis of such data might be indicative of a health problem and finally lead to the detection and treatment of a previously undiagnosed disease.

6. Subjects or carers trust the validity of the system too much

An issue which straddles both the topics of behavioral modification and regulatory validation (which will be discussed later) relates to the level of trust a user or carer has in the validity of a device. Certain risks may arise when either user or carer trusts a device excessively given how much it deserves to be trusted based on either its documented accuracy or how rigorously it was validated during regulatory approval. It is feasible that placing excessive trust in a device may ultimately place the user's health and wellbeing at more risk than if they did not use the device at all. A hypothetical scenario might involve an older relative living alone and using an automatic

fall detection pendant. The family may feel relieved of some burden of care, perhaps not checking on their loved one as regularly as before the device was used. If the detection sensitivity of the device is inadequate and a fall goes undetected, this could lead to a worse outcome for the user than if the family regularly checked on them with a phone call. This issue, that perhaps we are engineering new problems for ourselves by designing systems which encourage the substitution (rather than augmentation) of family care with technology, was previously raised by Redmond *et al.* [13]. Of course, similar scenarios may be envisaged for other wearable devices which monitor for physiological anomalies or adverse events.

Factors related to the social interaction between subjects using wearable sensors may also have a significant impact on the utility of such systems. If such sensors are visible to other people (which is the case, for example, with wrist-worn sensor devices), this may draw comments from an observer. This commentary may include personal and social feedback, which could be positive or negative, wanted or unwanted. Positive feedback by other people might relate to the attribution of positive personality characteristics, like self-control; negative feedback might emphasize the risks to the privacy of the bearer.

One way of anticipating and subsequently either preventing or leveraging such initially unintended effects would be to employ a user-centered design approach from the outset, using a participatory design process and an interdisciplinary team, including people with psychological skills, when developing and setting-up such wearable systems. A framework for doing this in a systematic way has already been proposed for behavior change programs based on serious gaming [14]. Once the system has been deployed, obtaining feedback from subjects on such particular aspects at regular intervals may be useful for assessing the existence and severity of unintended effects. For this, it may be necessary to design dedicated methods (e.g., questionnaires) to deal with this group of effects.

² In a personal communication a geriatrician pointed out that such situations repeatedly happen, for example, when the care situation changes and the patient goes from home care to institutional care. The same effect could also happen when a wearable adherence enforcement system is being deployed and taken too seriously by the subject.

Unintended Creation of Big Data and the Repercussions

As highlighted by Redmond *et al.* [13], there may be significant benefits to be gained by searching for associations within large datasets generated by users of wearable sensors worldwide. But allowing ourselves to be wholly or partly identified from our wearable sensor data may come at a cost; but is the price too high?

Insurance, Tax, and Lifestyle

Knowledge of health risks affect the premiums that health insurance companies offer to an individual. So with the rise of wearable sensors that monitor a myriad of medically valuable personal health parameters, insurance companies and employers are becoming interested in these technologies [15].

Many health issues, including comorbidities with obesity, are hugely affected by measurable or calculable lifestyle parameters, such as the intensity and frequency of physical activity undertaken in daily life. Knowledge of this information therefore has huge commercial value in the insurance world, as it can aid in further personalizing insurance premiums which may even change with time as a person's activity levels vary. Individuals who actively seek to improve and maintain their health may be rewarded with lower premiums, while sedentary individuals may be penalized with higher premiums. Employers may also encourage or request their employees to wear activity monitors, and reward healthy behavior. Furthermore, one could envisage taxes and levies imposed by a government seeking to incentivize a certain lifestyle for its citizens.

Benefits of Data Sharing

However, there is a very clear benefit to sharing de-identified data; including both wearable sensor data and other health, demographic, or lifestyle information. People around the world have begun to recognize the value of large-scale sharing; a concept often termed Open Data. Platforms for

sharing data already exist, such as PhysioNet [16] and The Health Data Initiative [17]. Large datasets are accessible through these platforms at no monetary cost to the user. Research fields that suffer from studies conducted on small datasets would benefit greatly from the pooling of data to create large sample sizes, and focusing intellectual labor.

For example, in geriatric fall prediction [18], pooling activity and falls data collected by wearable sensors across multiple studies, and making them open access, may enable the development of better prediction models. In fact, 'intellectual crowdsourcing' is the basis of Kaggle [19], an online platform where researchers and companies, both within and outside of healthcare, post datasets to crowdsource statistical and data mining labor via competitions to find the best predictive models.

Therefore, Open Data as applied to wearable sensors (and healthcare in general) may act as a means of harnessing two positive benefits: benefiting the individual with personalized medicine as they can benchmark themselves against the wider population; and benefitting research with larger datasets, greater study power, and fewer false positives.

The Great Debate

Whether the uses of health data obtained from wearable sensors as described above is fair and utilitarian, or is excessively intrusive or ultimately damaging to society, is still up for debate. In the case of the insurance companies mentioned above, the situation may not necessarily be so sinister: such a system may in turn encourage healthier lifestyles for the individual. If the long term health benefits of exercise are not sufficient incentive for someone to lead an active lifestyle, the immediate financial benefits arising from a wearable-sensor-enabled, personalized insurance policy may be.

In a survey, Atienza *et al.* [20] found that participant's views of their privacy and security (which in the context of insurance is related to their freedom of lifestyle) was dependent on what kind of information was being transmitted, who was accessing

the information, and where and when the information is being accessed. The survey uncovered a wide variety of attitudes among various diverse groups, but participants were generally willing to weigh up security/privacy with benefits, but the balance changed depending on how 'stigmatizing' the information was. It is interesting and perhaps rational that people are willing to surrender some privacy/safety if there are sufficient other benefits to be yielded from a monitoring system. Balancing this risk and reward will likely be a central to many future commercial, legislative, and regulatory developments in this area.

Regardless of which perspective we take, these issues must be confronted in the near future, which means that associated aspects, such as data storage, security, ownership, and visibility, must be swiftly and informedly addressed, and legislation must keep pace with these technological developments.

Privacy and Security

The previous sections have discussed the unintended consequences associated with wearable sensors, assuming the devices and data they generated are completely secure and accessed legally. This section will change tack somewhat to discuss the unintended security weaknesses of wearable sensors and how they expose user privacy and safety to varying degrees, depending on the nature of the security breach.

Networked medical devices, mHealth technologies and cloud services are a double-edged sword; they have the potential to play a transformational role in healthcare but may also be a vehicle to expose patients and healthcare providers to safety and cybersecurity risks such as being eavesdropped on, hacked, having their technology infected with malware, and being vulnerable to unauthorized access [21]. These actions are performed by adversaries with different motivations; for example, individuals or organizations that either act deliberately or maliciously (or work outside their ethical and legal remit), or are just unaware of the consequences of their actions.

Therefore, from a sensor perspective, the unintended consequences of wearable sensor use in healthcare can be summarized by two statements:

1. Wearable sensors are prone to the (accidental or deliberate) exposure of patient information and patient privacy;
2. There is a high level of trust in sensors and data provided by sensors, therefore they are susceptible to attacks that may potentially harm their users.

The following sections will examine these statements in more detail.

Privacy and Wearable Sensors

The National Committee for Vital and Health Statistics (NCVHS), a key advisory committee to the US Department of Health and Human Services, defines health information privacy as an *individual's right to control the acquisition, uses, or disclosures of his or her identifiable health data* [22]. The advent of wearable medical sensors requires us to broaden this definition, which leads to the following privacy categories [23]:

- **Device-existence privacy**
An unauthorized party should not be able to remotely determine that a patient has one or more sensors. For example, an adversary might be a potential employer willing to discriminate against the ill, a member of an organized-crime group seeking to sell a valuable device, or, in the case of military personnel, an enemy operative.
- **Device-type privacy**
If a device reveals its existence, its type should still only be disclosed to authorized entities. Patients might not wish to broadcast that they have a particular device for many reasons. For example, the device might treat a condition with a social stigma, it might be associated with a terminal condition, or it might be extremely expensive.
- **Specific-device ID privacy**
An adversary should not be able to identify or track individual sensors. This is analogous to the concern about the use of persistent identifiers in RFIDs, Bluetooth,

and 802.11 media access control (MAC) addresses to compromise an individual's location privacy.

- **Measurement and log privacy**
An unauthorized party should not be able to learn private information about the measurements or audit log data stored on the device. The adversary should also not be able to learn private information about ongoing telemetry.
- **Bearer privacy**
An adversary should not be able to exploit a sensor's properties to identify the bearer or extract private (non-measurement) information about the patient. Such information includes a patient's name, medical history, or detailed diagnoses.
- **Data integrity**
An adversary should not be able to tamper with past device measurements or log files or induce spurious modifications into future data. No one should be able to change when an event occurred, modify its physiological properties, or delete old events and insert new ones. A patient's name, diagnoses, and other data should be stored in a manner that is physically protected and tamper-proof.

Adversaries and Attack Mechanisms

Attacks on devices can be conducted by different types of adversaries with different motivations, different technical skills, and different sets of credentials:

- **Passive adversaries**
Such adversaries eavesdrop on signals transmitted by the sensors and network infrastructure (i.e., gateways);
- **Active adversaries**
These adversaries can interfere with legitimate communications and initiate malicious communications with sensors and network equipment. They are also capable of manipulating sensor hardware; e.g., performing on-board probing;
- **Coordinated adversaries**
Two or more adversaries might coordinate their activities; for example, one adversary would be near a patient and another near a network gateway;
- **Insiders**
These include healthcare professionals,

software developers, hardware engineers, and, in some cases, patients themselves.

Adversaries can conduct attacks on devices by exploiting sensor and system weaknesses, including the following:

- Principal design flaws of hardware or software as well as insufficient physical protection (e.g., tamper-proof device enclosures) [24];
- Software implementation flaws like untested or defective software/firmware or hard-coded passwords or access vulnerabilities deliberately introduced by developers for debugging, testing and remote maintenance purposes;
- Configuration issues, including misconfigured networks or poor security practices [25];
- Maintenance issues, like the failure to install timely manufacturer security software updates and patches [21];
- User issues including the uncontrolled distribution of passwords;
- Open, easily accessible communication networks.

These weaknesses (in conjunction with the adversary's capabilities) are exploited using one or more sensor access mechanisms as listed in Table 1:

- *Direct invasive* and *direct non-invasive* access requires an active adversary to have direct physical access to a device.
- *Remote non-invasive active* and *remote non-invasive passive* attacks are conducted by adversaries over distance by exploiting the sensor's communication link.
- *Visual access* only requires an adversary, for example an insider, to have line of sight to the sensor.

The information that can be potentially collected or manipulated by an adversary falls under two categories:

- Sensor metadata; e.g., sensor type, sensor associations (multiple devices carried by single patient, i.e., a closed-loop insulin delivery system), sensor/patient location, sensor activity, and visibility patterns.
- Sensor on-board data; e.g., configuration parameters, sensor readings, sensor status data and patient information.

Sensor Connectivity

Wearable sensors provide a wireless RF interface that is either intermittently (e.g., pacemaker) or permanently (e.g., ECG monitor) enabled, in order to provide the exchange of data and or commands/configuration settings between the device itself and some base station [26]. The base station relays an end-to-end connection, typically via an IP network, between the device and a management station / data concentrator [27]. The latter can be either deployed locally (e.g., a hospital server) or remotely (e.g., a cloud-based service).

Depending on the underlying RF technology, a gateway can be a customized base station (e.g., ZigBee), an access point (e.g., Wi-Fi) or a cellular base station (e.g., GSM). Multiple geographically-distributed gateways form a mesh that allows roaming (e.g., the movement of a sensor) over an extended geographical area beyond the transmission range of the underlying RF technology (i.e., an extended service set in IEEE 802.11), while a high density of gateways provide either increased data throughput, redundancy, or even sensor localization via signal trilateration. Gateways are interconnected via a high-speed backbone infrastructure. Note that while many sensor network technologies support ad-hoc networks that form complex meshes, where packets are routed in a best-effort approach via multiple hops to reach the gateway, medical sensor networks tend to form a simple single hop star architecture with the gateway in the center.

All RF technologies use a common broadcast medium split into individual channels that is shared between all networked devices (i.e., the 2.4 GHz GSM-band used in IEEE 802.11-based networks), but differ with regard to transmission range and transmission bandwidth (see Figure 1), which in turn determine their energy requirements.

The data to be exchanged between a wearable sensor and a gateway is embedded in individual frames or packets, which are transmitted to a gateway. A packet consists of a header and a payload section. An entire network packet (excluding header sections that can be modified during the routing of the packet) can be authenticated (via

Table 1 Device access methods and their impact on privacy.

Access Category	Access Opportunity	Access Mechanism	Potential Impact on
Direct invasive	Theft	In-situ (on-board) probing	- Device type privacy - Specific-device ID privacy - Measurement and log privacy - Bearer privacy
Direct non-invasive	Theft, Opportunistic	- Direct attachment (USB, serial) - User interface (keypad, LCD)	- Device type privacy - Specific-device ID privacy - Measurement and log privacy - Bearer privacy
Remote non-invasive passive	Opportunistic	(Wireless) network access; e.g., interception of data communication	- Device existence privacy - Device type privacy - Specific-device ID privacy - Measurement and log privacy - Bearer privacy
Remote non-invasive active	Opportunistic	(Wireless) network access; e.g., injection, interruption, replay and modification of data communication	- Specific-device ID privacy - Measurement and log privacy - Bearer privacy
Visual	Opportunistic	Line of sight	- Device existence privacy - Device type privacy - Specific-device ID privacy




Fig. 1 Data rates (bits per second) versus transmission range (meters) for RF technologies commonly used for wearable sensors.

a cryptographic checksum) to avoid flight manipulation, and its payload can be encrypted (to provide secrecy), but in any case the source and destination address of a packet that determine the origin of a packet and its destination is readable plaintext information. These addresses are for example static hardware MAC address assigned to the device during manufacturing.

Risks of Remote Non-invasive Passive Sensor Access

Wireless broadcast transmissions with the exemption of cellular networks can be easily passively eavesdropped by a sniffer, even over a distance if a high-gain antenna is used. Examples include USB-powered ZigBee sniffer devices, or Wi-Fi network interface cards configured in promiscuous mode. Captured and time-stamped network packets can be further examined using packet analyzer software, like, for example, the open source tool Wireshark. This kind of interaction, which is completely passive and undetectable allows the determination of:

1. What device (identified by the packet source address) is in vicinity of the sniffer at what time;
2. What the temporal data direction and data volume patterns for a given device are - even if the payload is encrypted;
3. Whether devices are stationary or move (derived from temporal variations of the received signal strength at the sniffer end), which allow to determine movement patterns of the user/patient.
4. The device type identification via temporal correlation of data transmission volumes (fingerprinting). Note that fingerprinting works even if packet payloads are encrypted.
5. Clusters of multiple devices (based on correlated signal strengths of different devices) that could be associated with a single user/patient.
6. What patient is associated with what device (and subsequently what device type), if the appearance (say, admission to hospital) of a user can be correlated to the detection of a sensor.

Therefore eavesdropping renders both device-type privacy as well as specific-device ID privacy (as defined above) useless, while device-type privacy and bearer privacy are potentially undermined. Concrete examples on how eavesdropping can be used to undermine different privacy aspects are outlined below.

Examples

- A network sniffer, strategically positioned in a hospital ward or step-down facility where patients carry wireless sensors, allows determination of all the above. For example, this information could be used to estimate utilization levels of the ward, patient activity/movement patterns, and average length of stay;
- Similarly, a sniffer positioned in a sports ground would allow determination of the presence of individuals/athletes that carry fitness trackers. This information, if collected in real-time, could be used to coordinate burglaries of these individuals' homes. Note that similar incidents were reported in the past, whereby criminals used information from social media websites to track the location of individuals;
- A health insurance provider that determines an association between a potential customer and a wireless sensor associated with a chronic disease (i.e., an insulin pump), could be alerted to a pre-existing condition of the customer;
- Petty criminals that target expensive medical sensors could actively listen for such devices via fingerprinting;
- Likewise, petty criminals that target vulnerable individuals (that are an easy prey once identified) could actively scan (via fingerprinting) for devices that indicate the user might suffer a debilitating disease.

Risks of Remote Non-invasive Active Sensor Access

Conventional cyberattacks on enterprise systems or critical infrastructure are the mainstream equivalent of this kind of access. Here an adversary attempts to gain

access to a sensor, for example by guessing or reverse-engineering credentials that allow him or her to access the system, or by actively looking for bugs in the software/firmware that can be exploited; for example, known issues with a sensor's web-server or protocol stack implementation. A typical, remotely conducted attack consists of up to five phases, namely: (i) reconnaissance, to understand the architecture and components of a sensor or sensor network; (ii) scanning, to look actively for vulnerabilities in the target system that can be subsequently exploited; (iii) gaining access; (iv) maintaining access, and; (v) covering tracks. The adversary can be based within the RF range of a wireless network itself, or can enter the network from a distance via its gateway.

A successful attack on a sensor or an entire network can have far reaching consequences with potentially huge impact on patient privacy and patient safety [28] as shown in the following examples. While such an attack requires an in-depth knowledge and understanding of the target system, it affects specific-device ID privacy, measurement and log privacy, as well as bearer privacy and patient safety.

Examples

- In 2011, a security researcher showed how an insulin pump manufactured by a well-known medical device company could be remotely attacked and its entire insulin supply administered to the patient, therefore causing a possible lethal insulin shock [29];
- Similarly, sensor readings can be systematically manipulated, for example low SpO₂ readings of a hypoxemic patient being increased to a normal 95-100% saturation level via direct manipulation of firmware or the upload and injection of malware into the device;
- Even apparently secure device communication setups can be compromised if faulty firmware is used. A well-known software defect is the Heartbleed bug [30] of the cryptolibrary OpenSSL, which allowed adversaries to retrieve key material used to authenticate and encrypt device communication within the TLS (Transport Layer Security) protocol.

Risks of Direct Invasive and Direct Non-invasive Sensor Access

Direct actions are far more tangible and generally better understood than the remote access actions above. Both approaches require the adversary to have physical access (either via theft which will be eventually discovered, or by the opportunistic temporarily access of a device at its momentary location).

Direct invasive access via in-situ (on-board) probing is a “brute-force” attack mechanism which has been successfully applied to other types of connected embedded devices; i.e., smart meters [31] or phones [32].

Invasive access involves the probing of sensor hardware with the aim of retrieving data (e.g., firmware, system settings, patient data, recordings or credentials) from it. This data is stored on volatile (e.g., DRAM or SRAM) or non-volatile (e.g., Flash) memory. Hardware access can be accomplished by different means, including direct pin-probing of memory chips and CPU, access to local system buses (i.e., I2C or SPI) that connect the CPU with other system components, or via the JTAG interface, that provides direct access to all CPU subsystems including general purpose registers, special function registers and internal RAM or Flash memory. The most sophisticated attacks place probes directly onto the silicon and provide gate-level access to a chip.

On-board sensor data is generally stored as unencrypted plaintext or binary data. Once it has been accessed and extracted, data can be further analyzed and dissected via off-the-shelf editing tools. Depending on the extent and content of such a memory dump, device type privacy, specific-device ID privacy, measurement and log privacy and bearer privacy will be compromised.

Non-invasive direct access does not require direct access to a sensor’s hardware, but uses instead the device’s interface (i.e., display/keypad or USB) to extract device data; for example, a sensor configuration menu. Here an attacker uses a legitimate access path, but exploits a sensor design or configuration flaw, whereby access to the aforementioned menu is either not sufficiently (password) protected, or uses credentials

(e.g., passwords) that are known to the attacker or can be guessed via a brute force approach (i.e., a 4-digit pin code).

Here, device type privacy, specific-device ID privacy, measurement and log privacy as well as bearer privacy can be compromised.

Examples

- Incorrectly decommissioned sensors can be acquired by technically competent adversaries and examined via in-situ probing. Any patient-related information (i.e., personal or insurance details) that have been loaded into the device when it was assigned to the patient, but not properly erased afterwards or prior to decommissioning, can be potentially retrieved;
- In 2014, security researchers in the US demonstrated how a network enabled infusion pump could be remotely manipulated via its network interface. The culprit here was a default factory password that was never changed;
- In a similar attack the researchers manipulated a Bluetooth-enabled defibrillator to deliver random shocks;
- A patient that can access and (without being aware of their wrongdoing) change sensor settings (i.e., the output rate of a wearable insulin pump) via the device’s user interface unknowingly put themselves at risk.

Risks of Visual Sensor Access

Device existence privacy, device type privacy and specific-device ID privacy may be compromised if an adversary has visual (line of sight) access to a wearable sensor and its carrier. While the sensor hardware including the RF interface follows the general miniaturization trend of silicon design (and can be consequently hidden under the carriers cloth), a sensor’s actuators are restricted in terms of size and location. For example:

- ECG electrodes have a minimum diameter for better connectivity to the patient’s skin, even though they can be easily hidden under clothes (i.e., chest electrodes of a Holter ECG).
- Non-invasive blood pressure monitoring

systems (based both on the auscultatory and oscillometric method) require an inflatable cuff wrapped around the upper arm or wrist that reduces mobility, particularly when the cuff is being inflated. Also, the sensor’s compressor produces an audible humming sound;

- SpO₂ sensors are attached to parts of the limbs with good arterial visibility/translucency (i.e., an index finger) and can be easily identified;
- An insulin pump requires a small reservoir with a capacity of typically 300-400 IUs, in addition to the pump itself. It is often attached to a user’s belt, making it visible to bystanders.

Examples

- Petty criminals that target expensive medical sensors could actively monitor for such devices attached to their victims;
- Likewise, petty criminals that target vulnerable individuals can look out for victims that carry devices which indicate a debilitating disease;
- Patients could be stigmatized by their environment if the device they are visibly carrying links them to an illness or disability (i.e., a person with type 1 diabetes wearing an insulin pump on their belt). The psychosocial implications of such effects were discussed in an earlier section.

Security of Data in the Cloud

Cloud computing and storage solutions give healthcare providers the capability to store and process patient and sensor data in third-party data centers, which are made available in a variety of service models (e.g., software-, platform-, or infrastructure-as-a-service) and deployment models (e.g., private, public, hybrid, and community cloud).

However, cloud services are prone to the same type of adversaries and to similar attack mechanisms as outlined before, which in turn can compromise measurement and log privacy, bearer privacy, data integrity, and consequently patient safety. The information that can be potentially collected or manip-

ulated by an adversary falls predominantly under the previously-defined category of on-board sensor data, and to a lesser extent under the sensor metadata category.

The outsourcing of data storage and data processing to the cloud makes it necessary to address privacy and data security aspects both from a cloud provider and a client perspective; e.g., the provider must ensure that their infrastructure is secure and that their clients' data and applications are protected, while the client must take measures to restrict access to its data and services stored in the cloud.

Failing to implement such measures properly has huge implications as highlighted by a recent report [33]:

- Insecure cloud interfaces and APIs, including anonymous access and/or reusable tokens or passwords, clear-text authentication or transmission of content, as well as unpatched cloud software, allows an external active or passive adversary to gain access to patient and sensor data stored in the cloud;
- Likewise, omitted background checks for employees who have physical access to the servers in a data center, as well as a lack of monitoring for suspicious activities within a data center, enables an insider adversary to gain access to patient and sensor data stored in the data center;
- Improperly managed access control mechanisms (i.e., weak passwords) can lead to account or service hijacking by an adversary.

Various international working groups and organizations including the Cloud Security Alliance, currently address all these issues by promoting best practices for providing security assurance within cloud computing.

Examples

- A step-down unit of a private hospital, providing intermediate patient care for those moving between an intensive care unit and a normally-staffed inpatient division, is equipped with a ZigBee-powered wireless sensor network for continuous patient monitoring. The sensors are connected to a number of wireless gateways located in the step-down unit,

which in turn upload all sensor data to a cloud-based data center. A bug in the data center's protocol stack software (the aforementioned Heartbleed bug) is exploited by an external attacker, who subsequently gains access to all sensor data stored in the cloud.

- In a similar scenario a disgruntled employee of the data center, who has access to the sensor data stored on the centers' servers, systematically manipulates sensor readings of patients.
- Again in a similar scenario, the hospital's IT staff uses an insecure protocol parameter (the so-called IKE aggressive mode used by the IPSec protocol) to setup an apparently secure tunnel connection between the hospital and the data center. An external attacker exploits this vulnerability, breaks into the connection and is able to eavesdrop on all sensor and data communication between the hospital and the data center.

Regulation and Patient Safety

This section discusses the reduction or elimination of some of the unintended consequences associated with the use of wearable sensors, as outlined above, and highlights a number of challenges facing regulatory bodies charged with this overseeing this task.

The Current State of Wearable Sensors Regulation

The FDA has said it will not regulate wearable sensors designed purely for lifestyle purposes, such as those that generally promote health and fitness [34]. The FDA defines general wellness devices as, "...products that meet the following two factors: (1) are intended for only general wellness use, as defined in this guidance, and (2) present a very low risk to users' safety" [34]. Examples include Fitbit and numerous other apps, like MyFitnessPal.

However, the rate of development of medical apps is expected to increase over

the next few years [35]. Torous *et al.* (2016) of APA's Smartphone App Evaluation Task Force have noted that there are already some 165,000 healthcare apps directly available to patients and clinicians [36]. Medical smartphone apps, but not the platform they run on (e.g., iPhone) are currently classified as "medical devices" by the FDA and treated in exactly the same way as a piece of medical equipment [37, 38]. This is an interesting stance which may be extended to other devices/platforms, where only the software is considered to be a medical device, but the hardware is regulated at the same level.

The connectivity of wearable sensors complicates the picture further though. As discussed earlier in the paper, many wearable sensors (and smartphones) interact with cloud computing and storage resources. This interaction extends the challenge of regulating wearable sensors far beyond the device and its software, and leads to the requirement for a medical-grade cloud. Regulations have yet fully catch up with the cloud revolution and it remains unclear how cloud security requirements should be validated, although penetration testing is one possible engineering solution which has been proposed [39]. The rise of the IoT and personal connected healthcare are likely to increase complexity even more, where the tools used to manage one's health become increasingly distributed throughout the environment.

It is clear, given the ever-rising complexity and interconnectivity of healthcare technologies, it is infeasible for regulatory bodies to take sole responsibility for protecting patient safety and privacy from unintended vulnerability.

Who is Responsible for Preventing Unintended Consequences?

In the clinical healthcare regulatory sphere, a pertinent question to ask is who is responsible for mitigating the risk of unwanted consequences, such as security breaches. The responsibility of medical-grade wearable sensors would seem to be the responsibility of the FDA and other regulatory bodies who are charged with overseeing their certification.

In 2009, the FDA published a reminder that security issues in healthcare was a shared responsibility between medical device manufacturers and the facilities (e.g., hospitals) that use them [40]. Recently, the FDA has said that security of connected medical devices is an IT problem for the IT industry to solve.

Part of the challenge of security in the regulated healthcare industry lie with the fact that medical device manufacturers and the healthcare industry in general do not have a sufficient grasp of the technical concepts of IT and security, and IT personnel are somewhat disconnected from the regulatory mindset. Hayhurst [41] suggests that security must be a collaborative effort between clinical and biomedical engineers, all of whom must have a basic understanding of the issues (IT security concepts, etc.). It is arguable that the responsibility is shared between all three parties: the IT industry; the medical industry; and healthcare facilities. Certainly, the economic burden of regulating all medical devices and apps is a strong motivation for the FDA to place some responsibility on the IT industry.

The FDA's Reasoning for Shared Responsibility

The risk of negative security consequences in healthcare is becoming an increasing concern for regulatory bodies. With infinite possibilities for new devices and medical apps, how is it possible for regulatory bodies to rigorously assess them all? There are guidelines in place, along with regular publications discussing the FDA's current thinking on the issues, which are important for manufacturers and IT personnel. Ensuring device quality, safety, and consistency to predetermined specifications is the core focus of the regulatory bodies. The FDA recognizes that a network-connected device, such as a medical-grade wearable sensor, can pose a risk to patients, and they traditionally require that manufacturers submit a body of evidence justifying their design decisions, based on risk analysis and sound science, including validation data, using the "least burdensome approach" [42].

Design Control as a Framework to Eliminate Unintended Consequences

Manufacturers and third party developers who want to market wearable sensors for medical and clinical purposes in the USA must adhere to the FDA's Quality System Regulation (QSR), which has requirements for every aspect of the design and development of the wearables. Each country has an authority that regulates medical devices and products, but the FDA is a good representative example. QSR is also known by its code name 21 CFR 820.

Design Control is a subpart of QSR, which includes requirements for design validation of medical devices, such as medical grade wearables. The FDA says, "Design Controls are based upon quality assurance and engineering principles", and has issued a guidance document for manufacturers [43].

Figure 2 illustrates where design validation fits into the overall Design Control subsection of FDA CFR 820 [43], which is necessary for medical grade wearable sensor developers to implement, document and submit as part of their marketing application dossier. This is explicitly defined by the FDA.

Validation as an Important Part of Design Control

As illustrated in Figure 2, validation is a main step in the Waterfall Development Model issued by the FDA. Design Validation is an important concept for wearable sensor developers who wish to market them to the clinical community. It is defined by the FDA as, "...confirmation by examination and provision of objective evidence that the particular requirements for a specific intended use can be consistently fulfilled. Design conforms to user needs and intended uses".

The FDA's current thinking on quality-related issues is reflected in regular updates and supplementary guidelines. The reason medical device network security is more important than other types of security is that lives depend on it.

Rakitin [44] presents an interesting explanation of design validation and risk management which could be applied to wearable sensors: "Design validation and risk management are examples of required

activities performed by medical device manufacturers to help ensure that devices are as safe as 'reasonably practical.' While design validation and risk management have improved the safety of medical devices, the effectiveness of these activities is directly related to the ability of the device manufacturer to understand and simulate the disparate networking environments within which these medical devices are used. Increasing the effectiveness of design validation and risk management in complex networking environments will require the full cooperation and active participation of stakeholders, including medical device manufacturers, IT network equipment suppliers, clinical and biomedical engineers, and IT staff, as well as regulators... Among the many challenges facing device manufacturers is performing design validation '...under actual or simulated use conditions.' When 'actual use conditions' include connecting medical devices to a healthcare organization's network, device manufacturers must somehow address the fact that every healthcare organization's network is different. The safety and efficacy of medical devices can often be affected by the network the device is connected to. Some device manufacturers perform design validation activities with their medical devices connected to the healthcare organization's network."

Security of open wearable sensor devices is clearly paramount, but are there currently any specific validation techniques that can ensure there will be no unintended consequences, with 100% accuracy? FDA has purportedly looked at security requirements validation techniques. "Fuzzing" is one technique that was researched by the FDA [45]. This is a type of penetration testing, or ethical hacking. This technique, as the name suggests, uses a fuzzy approach of bombarding the device (e.g., wearable sensor) with almost correct, but slightly incorrect IP packets to check if they cause it to behave incorrectly. In that case, fuzzing would be an example of a validation tool.

This research by the FDA would be highly significant, but the FDA has never officially published any information about specific tools or techniques to validate medical devices or wearable sensors,

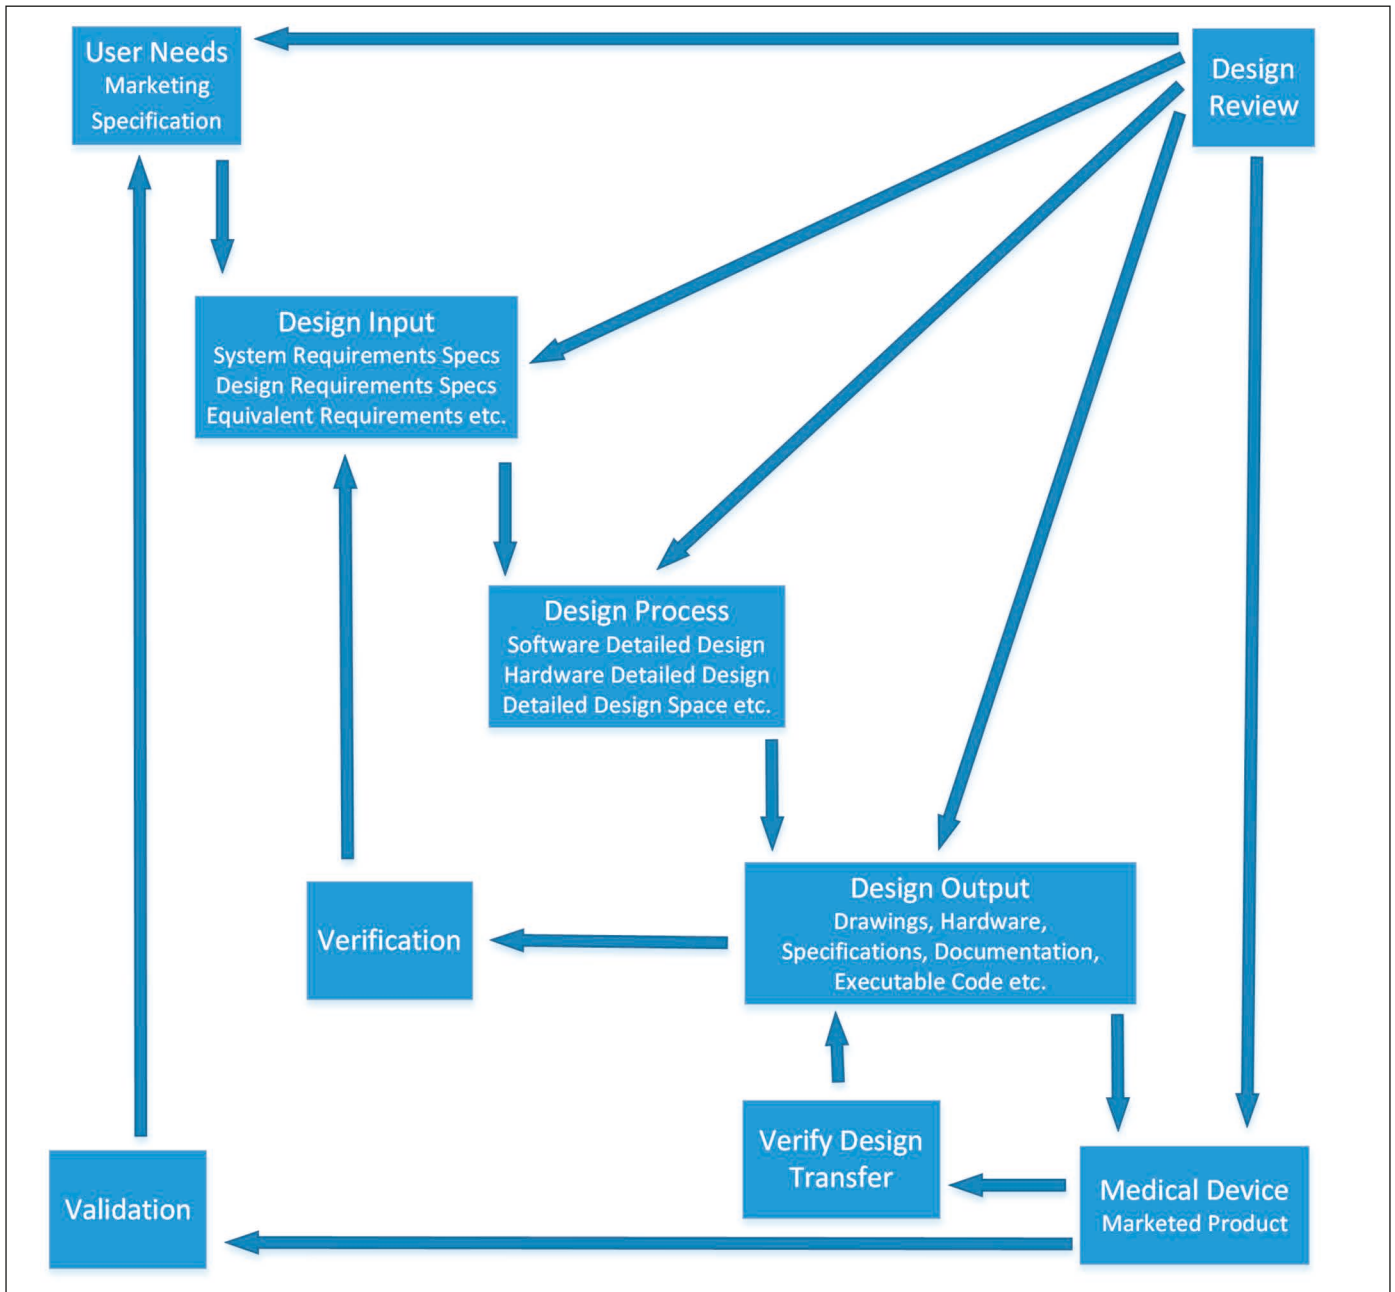


Fig. 2 Application of Design Controls to Waterfall Design Process. Documented evidence of following this process is a regulatory requirement for device manufacturers who wish to market their devices in the USA. It aims to prevent unintended consequences impacting negatively on patients and ensure that wearables (and other medical devices) consistently meet predetermined specifications [43].

nor officially expressed interest in such specificity. However, it is useful to know that the level of effectiveness of particular validation tools is a factor in preventing unintended consequences, and the FDA clearly realizes this. Indeed, it was not until early 2016 that the FDA published

solid guidance on device security, which states that pre-market validation of devices is not enough for security purposes, and validation must continue throughout the post-market life of the device until retirement [46]. This much-needed recommendation on security came in the wake of a

high profile recall of an approved medical device by the FDA on security grounds. It seems the security of wearable sensors and the ecosystem that supports them cannot be guaranteed; there are too many variables and unknowns. Security breaches can only be monitored for, not fully prevented.

Is Regulation Restricting Innovation for Wearable Sensors?

FDA approval is a potentially marketable feature for wearable sensors. Julie Papanek of the venture capital firm Canaan Partners, who invests in wearables start-ups, has said, “Consumers, doctors, payers all want to know if a product provides a clinical benefit... Working with the FDA is the one way to get the ability to market that benefit” [47]. However, obtaining medical devices certification comes with some disadvantages too.

Of course, adhering to regulations is a requirement for marketing a medical device to the public, but this requirement may inhibit innovation. In an interview with The Telegraph in early 2016, Apple CEO Tim Cook commented that they would not put the Apple Watch through the FDA process as adhering to regulatory requirements into the future would prevent them from innovating the product in future, to the degree needed to stay competitive in the marketplace [48].

Related to this, there is an additional disadvantage of regulation that could hinder innovations designed to improve wearable sensor safety, in addition to general marketability. What about the case of occasional software updates and patches that are required to improve security – must the wearable sensor go through the long re-validation process every time a software update is required?

In 2014, FDA clarified earlier guidelines, requiring that any plans for software patches or updates be included in the original filing of the marketing application: “The guidance recommends that manufacturers submit their plans for providing patches and updates to operating systems and medical software”. However, there are other cases where the lengthy and costly re-validation process may stifle innovation.

Numerous studies have looked at techniques to improve security. One popular technique mentioned in the literature is to use a jamming signal [49] to obfuscate the communication between a wirelessly connected device and its associated equipment (e.g., diagnostic or controlling equipment). Technology like this may become increasingly relevant to wearable sensors in the IoT. The signal would be jammed at all locations,

but an antidote signal would unjam it at the receiving end. Are techniques like this feasible to be retrofitted to wearable sensor systems *already* on the market? If not, is that part of the problem, given the unlikelihood that any regulatory body would allow innovations like this to be implemented without re-validation?

Further to having a good grasp of the effectiveness of validation tools, the literature is abound with recommended best practices for implementing security measures, which have applications in wearable sensor design and development. Some of the security design recommendations for devices used in clinical settings [50] include:

1. Security should be addressed from the beginning, i.e., from the design phases;
2. Do not rely on secret algorithms or hiding the hardware (“security through obscurity”), and instead implement well-established cyphers;
3. Use well-established tools for verifying source-code during the design phase;
4. Use authentication for third party devices that communicate with the device (or sensor as the case may be).

Thus, it is clear that discussing consequences of wearable sensors cannot be divorced from the regulatory considerations that govern them. Academics that

study and write on the topic must have an understanding of both. Collaboration between both fields is necessary.

The IT Industry Responds to the Challenges

In order to meet or exceed regulations and build a body of supporting evidence for marketing applications, the IT industry has also come up with its own frameworks and accepted standards for managing risk. The Regulated Software Research Centre & Lero, The Irish Software Research Centre, published its presentation, “A Security Risk Management Framework for Networked Medical Devices” [51]. The group proposes using assurance cases and product risk analysis. The presentation uses the term assurance cases to describe “a body of evidence organized into an argument demonstrating some claim that a system holds i.e. is acceptably safe” [51]. Assurance cases are recommended for demonstrating safety, security, or reliability of a system. Specifically, they give a general structure as follows:

The following steps are proposed [51]:

1. Must make a claim or set of claims about a property of a system;

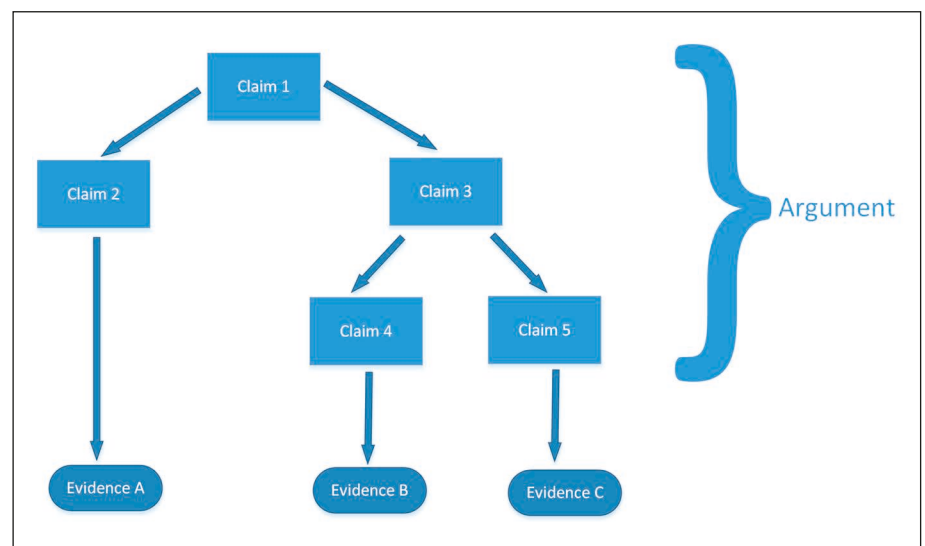


Fig. 3 General Assurance Case structure for demonstrating safety, security, or reliability of a system [51].

2. Provide a set of arguments;
3. Make clear the assumptions and judgments underlying the arguments;
4. Produce the supportive evidence.

Increased use of software on devices as well as device communication abilities are cited as a problem background to the framework. This will surely be a key interest to wearable sensor developers. The group proposes the IEC/TR 80001-2-8 standard incorporates FDA and other guidance [51].

Conclusions

This paper has discussed how the enormous increase in the use of wearable sensor technologies for disease, wellness, and lifestyle monitoring, may result in several unintended consequences. These unintended consequences relate to advantages or unwanted: modifications of user behavior; uses of information; security vulnerabilities; and regulatory challenges.

It is clear that many unintended modifications in behavior can be mitigated or eliminated through more intelligent and inclusive design processes. It is perhaps expected that as new wearable sensors and their associated software ecosystems are developed, we will see new and interesting perversions of behavior relative to what was intended by design.

On the creation of big data sets from wearable sensors and other sources, legislating who has access to that information (government or an insurer) and whether it can be used to discriminate against an individual, will likely be the source of much legal and political wrangling in the years to come, as we attempted to get the greatest good from wearable sensors, but at the right cost to our privacy.

We have also discussed how the wireless and personal nature of wearable sensors exposes them to a number of vulnerabilities which might expose private items of information. Again, a risk versus benefit trade-off is at play here. We want devices with long battery lives and small sizes, but such devices are more limited in how well they can secure information. Similarly,

we must consider how extreme the consequences are if a security breach occurs. Certainly, devices which have an actuator (such as a defibrillator, or insulin pump) should require the greatest level of security, whereas, sensor-only devices may be relatively less secure.

While we see regulation plays a role here in guiding the medical device industry, the burden of validating the function and security of the medical device market is becoming infeasible for regulators. With the rapid proliferation and advancement of mobile phone and wearable technology, and the rate at which software apps are being developed to interpret these physiological data collected, it will require a collaborative effort between device manufacturers, regulators, and end-users to strike the right balance between the risk of unintended consequences occurring and the incredible benefit that wearable sensors promise to bring to the world.

References

1. International Data Corporation. (27 Aug 2015). Apple Debuts at the Number Two Spot as the Worldwide Wearables Market Grows 223.2% in 2Q15, Says IDC. Available: <http://www.idc.com/getdoc.jsp?containerId=prUS25872215>
2. International Telecommunication Union. ITU Internet Reports: The Internet of Things. Geneva; 2005.
3. Schreier G. The Internet of Things for Personalized Health. *Stud Health Technol Inform* 2014;200:22-31.
4. Jara AJ, Zamora-Izquierdo MA, Skarmeta AF. Interconnection Framework for mHealth and Remote Monitoring Based on the Internet of Things. *IEEE Journal on Selected Areas in Communications* 2013(31):47-65.
5. Dohr A, Modre-Osprian R, Drobnic M, Hayn D, Schreier G. The Internet of Things for Ambient Assisted Living. Proceedings of the ITNG 2010 - 7th International Conference on Information Technology : New Generations. Las Vegas, Nevada, USA; 2010. p. 804-9.
6. Zhang J, Song Y-L, Bai C-X. MIOTIC study: a prospective, multicenter, randomized study to evaluate the long-term efficacy of mobile phone-based Internet of Things in the management of patients with stable COPD. *Int J Chron Obstruct Pulmon Dis* 2013;8:433-8.
7. Hafezi H, Robertson TL, Moon GD, Au-Yeung K-Y, Zdeblick MJ, Savage GM. An Ingestible Sensor for Measuring Medication Adherence. *IEEE Trans Biomed Eng* 2015(62): 99-109.
8. Stone M. (2014, accessed 08/01/2016). Smart-phone Addiction Now Has A Clinical Name. *Business Insider Australia*. Available: <http://www.businessinsider.com.au/what-is-nomophobia-2014-7>
9. Deb A. Phantom vibration and phantom ringing among mobile phone users: A systematic review of literature. *Asia Pac Psychiatry* 2015 Sep;7(3):231-9.
10. Smith MW. A Fitbit fanatic's cry for help: I'm addicted to steps. Lessons from a year spent fully quantified. *Washington Post*; 2015. Available: <http://www.washingtonpost.com/news/to-your-health/wp/2015/05/11/a-fitbit-fanatics-cry-for-help/>
11. Stegemann S, Baeyens JP, Cerreta F, Chanie E, Löfgren A, Maio M, et al. Adherence measurement systems and technology for medications in older patient populations. *European Geriatric Medicine* 2012(3):254-60.
12. Shull PB, Jirattigalachote W, Hunt MA, Cutkosky MR, Delp SL. Quantified self and human movement: A review on the clinical impact of wearable sensing and feedback for gait analysis and intervention. *Gait & Posture* 2014(40):11-9.
13. Redmond SJ, Lovell NH, Yang GZ, Horsch A, Lukowicz P, Murrugarra L, et al. What Does Big Data Mean for Wearable Sensor Systems? *Yearb Med Inform* 2014(9):135-42.
14. Cheek C, Fleming T, Lucassen MFG, Bridgman H, Stasiak K, Shepherd M, et al. Integrating Health Behavior Theory and Design Elements in Serious Games. *JMIR Mental Health* 2015(2):e11.
15. Olson P. Wearable Tech Is Plugging Into Health Insurance. *Forbes Tech*; 2014. Available: <http://www.forbes.com/sites/parmyolson/2014/06/19/wearable-tech-health-insurance/>
16. PhysioNet. (Jan 5, 2016). Available: <http://www.physionet.org/>
17. HHS Idea Lab. (Jan 5, 2016). Available: <http://www.hhs.gov/idealab/what-we-do/health-data/>
18. Shany T, Wang K, Liu Y, Lovell NH, Redmond SJ. Review: Are we stumbling in our quest to find the best predictor? Over-optimism in sensor-based models for predicting falls in older adults. *Health Technol Lett* 2015 Aug 3;2(4):79-88.
19. Kaggle. (2016, accessed 05/01/2016). Available: <http://www.kaggle.com/>
20. Atienza AA, Zarcadoolas C, Vaughn W, Hughes P, Patel V, Chou W-YS, et al. Consumer Attitudes and Perceptions on mHealth Privacy and Security: Findings From a Mixed-Methods Study. *J Health Commun* 2015;20(6):673-9.
21. Deloitte. Networked medical device cybersecurity and patient safety. White Paper; 2013.
22. Kotz D. A threat taxonomy for mHealth privacy. Presented at the 2011 Third International Conference on Communication Systems and Networks; 2011.
23. Halperin D, Heydt-Benjamin TS, Fu K, Kohno T, Maisei WH. Security and Privacy for Implantable Medical Devices. *IEEE Pervasive Comput* 2008(7):30-9.
24. Al Ameen M, Liu J, Kwak K. Security and Privacy Issues in Wireless Sensor Networks for Healthcare Applications. *J Med Syst* 2010(36):93-101.
25. Symantec. Networked Medical Devices: Security and Privacy Threats. White Paper; 2012.

