ICTs and the Health Sector:
Towards Smarter Health and Wellness Models

OECD, October 2013

Chapter 9

Building a smarter health and wellness future: Privacy and security challenges

Carl A. Gunter

This report explores emerging privacy and security challenges for health information technology (HIT) that call for new ideas. Six key challenges are identified and discussed: access controls and audit, trusted base, automated policy, mobile health, identification and authentication, and data segmentation and de-identification.

Carl A. Gunter, Building a smarter health and wellness future: Privacy and security challenges. Chapter 9 in ICTs and the Health Sector: Towards Smarter Health and Wellness Models, OECD, October 2013, pages 141-157.

Introduction

Advances in information technology (IT) promise to improve health and wellness by holding and managing detailed and precise records related to diagnoses and treatments and encouraging good lifestyle choices through applications like fitness monitoring. These benefits will come through improved abilities to collect, manage, share, and act upon digital information using computers and digital network links. Many of these technologies will use broadcast wireless communications, global network connectivity provided by the Internet, and shared hosting of data and computations based on cloud computing. Each of these technologies, and many others likely to be used, raise essential questions about the privacy and security of the data they store or transmit.

For example broadcast wireless data is easily "sniffed" so that even encrypted links can leak information through traffic analysis. The Internet suffers notorious problems with skilled and ethically challenged hackers who have access to systems across the world, and cloud computing is in its early days and displays shifting conventions about how personal data will be mined for the commercial beneft of the cloud provider.

This chapter explores emerging privacy and security challenges for health information technology (HIT) that call for new ideas. While there is much to be gained in security of HIT by simply applying procedures and protocols that have worked in other areas like the financial services sector, there are many special characteristics of HIT and trends in HIT that call for innovation. This may be either in the way existing techniques are applied or in the need for new techniques.

To see this in an example, consider the analogy between personal health records (PHRs) and personal online banking. PHRs make health care provider data about a patient available to the patient, just as personal banking makes bank data about a customer available to the bank customer. Online personal banking and PHRs thus have many privacy and security issues in common. There is a need for good authentication protocols (keys and passwords) and support for an encrypted communication channel. Patients, like banking customers, may want to merge information from multiple sources to get a unified view, as some financial services packages (like tax preparation software) enable. There is common need to share data with third parties: just as a patient needs to show medical records to a new doctor (like a specialist), a banking customer may need to show data to a financial entity (like a mortgage lender). However, beyond these similarities there are also critical differences. For instance, consumer financial data are relatively simple compared to medical data, which use a large and changing vocabulary of terms and codes for medical conditions and treatments that can be difficult for doctors to understand, let alone patients. This has consequences for privacy because patients need help understanding how to share their medical data with parties that could benefit from having it. This includes sharing with members of the medical profession of course, but also sharing with third parties like online vendors who offer to host medical data and provide analytic services. Even in circumstances where financial data is as complex as health data, the means and motivations for sharing are quite different. In addition, it is not the case that the financial services sector has solved its own security problems fully. Problems like identity theft remain a major challenge for financial services even as they are also becoming one for health care.

We consider six key areas where research is needed to improve techniques for privacy and security of HIT. Before attending to these key areas, it will first be worthwhile to provide some background on HIT policy developments that raise privacy and security issues, and to make some introductory remarks about the concepts of privacy and security. Then, we will concentrate on the following key areas: access controls and audit, trusted base, automated policy, mobile health, identification and authentication, and data segmentation and de-identification. We conclude with discussion of some cross-cutting concerns.

HIT directions driving privacy and security issues

There are a variety of important trends in HIT that drive issues with privacy and security. We begin by discussing a few of these. The concept of the learning health care system, which we describe in a moment, provides a framework for thinking about trends in health data use. Two specific areas illustrate the opportunity for learning based on HIT. First, health information exchange between providers, patients, researchers, and public health make provider data available for learning and, second, mobile health enables new types of data to be collected from individuals by monitoring information about their lifestyles.

Learning Health System (LHS) is an agenda developed by the Institute of Medicine (IOM) Roundtable on Value & Science-Driven Health Care (Grossmann et al., 2011): “Our vision is for a health care system that draws on the best evidence to provide the care most appropriate to each patient, emphasizes prevention and health promotion, delivers the most value, adds to learning throughout the delivery of care, and leads to improvements in the nation's health.”

Although this IOM statement is aimed at action in the United States, international efforts have similar goals. The stated goal for LHS is to reach a point at which clinical decisions are supported by accurate, timely, up-to-date information that reflects the best available evidence.

The connection of this agenda to HIT is the enabling capability of computers and digital networks to collect and transmit data that can be used to develop evidence and assure that the right information based on this evidence is in the hands of the provider or individual at the point a health decision is required.

In particular, HIT is a critical enabler of Health Information Exchange (HIE), a phrase that typically refers to the exchange of health data between diverse parties for the better care of individual patients. A paradigmatic example is enabling the primary care giver for a patient to send the patient's record to another provider where the patient needs specialized or emergency care. Such exchanges save money by avoiding unnecessary tests and can save lives by reporting safety critical information like medications and allergies.

Used as a noun, “HIE” is a system that facilitates exchange, often by setting up secure and standardized communications between providers to serve as an infrastructure for exchange. One key issue addressed by HIEs is interoperability. A typical architecture for such an HIE appears in Figure 9.1. The Electronic Medical Record (EMR) of a provider is (typically) a proprietary system that holds health records in a computer. A portal allows records from the EMR to be converted to a standard EHR format such as a Continuing Care Document (CCD). When a provider seeks a record for a patient from the HIE, the system returns a consolidation of such records obtained from the participating providers that have records on the patient. This consolidated record is then made available to a client user like a doctor in an emergency room that is treating the patient.

Figure 9.1. HIE architectures

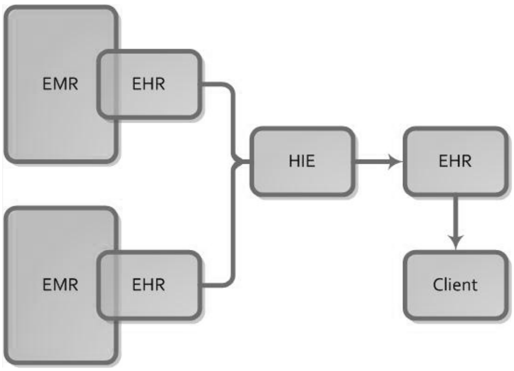

Source: Author.

There are many variations on this theme. For instance, the records can be pushed to the HIE central repository and stored there for quick recall, or they can be retrieved only on demand, so that records not requested are never seen by the HIE. These and similar choices affect a second key issue addressed by HIEs, namely privacy and security. We discuss these issues in some of the challenges below.

Mobile health (mHealth) concerns the increasing use of mobile sensors to monitor health and wellness conditions and, for certain types of devices, carry out actions to improve health (actuators). On the one hand mHealth devices like pedometers in cell phones assist with measurements that encourage healthy lifestyles. On the other hand, mobile implants, which can be broadly classified as mHealth devices, collect health vitals and can intervene with actuation in an emergency. For example, an implanted defibrillator intervenes with an electric shock when it detects a cardiac emergency.

Risks are quite different for these two extremes, but there are an increasingly large number of devices that fall between these extremes. There is interest in having insulin pumps that communicate with cell phones to facilitate viewing and sharing of data. Figure 9.2 illustrates a common pipeline of communications for mHealth devices. A sensor on (or within) the body collects data that flows to an mHealth device. This device talks, for instance, to a cell phone or wireless station (gateway device), which, in turn, sends the data over the Internet to an EHR, PHR, or EMR where it is available for viewing by an analyst (a doctor for instance). The pipeline can then be used in reverse for configuration or actuation. The overall pipeline can also be”short circuited”at various stages. For instance, the individual wearing the device might process and view the data on his or her device without sending it over the Internet.

Figure 9.2. mHealth pipeline

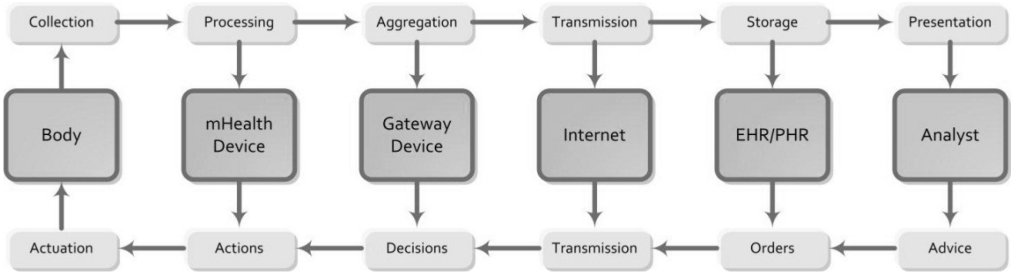

Source: Author.

Privacy and security

A few background comments on privacy and security will be useful for this discussion. First of all, it is helpful to distinguish between these two concepts. Privacy is notoriously difficult to define precisely (Nissenbaum, 2011), but it commonly refers to the desires and expectations of individuals about how, when, what, and to whom information about themselves is revealed to others. Privacy violations are violations of these desires and expectations. If Alice tells her friend Bob a personal fact about herself in confidence, like, say, that she has been diagnosed with cancer, and then Bob mentions this in a Facebook comment, Alice may well consider this a violation of her privacy. As recognized as far back as the Hippocratic Oath, it is essential to assure that patient privacy is respected during the provision of health care and treatment services. Without this assurance many individuals will not seek the health services they need.

By contrast, security typically refers to instances in which information is deliberately accessed and used for unauthorized or illegal purposes. For instance, if Alice responds to a phishing email and reveals her personal banking password to a hacker who accesses her bank account, then Alice is a victim of a security breach. Privacy and security are closely related. For instance, the attacker who accessed Alice's bank account will probably learn how much money she has, a fact she would not have revealed to a stranger.

The improper use of HIT can increase the danger of compromising the security and privacy of individuals. HIT allows health information to flow easily from one place to another, a property sometimes called "liquidity". Providers exploit liquidity to share data with payers, researchers, public health, and other providers. These flows may well violate expectations of subjects. For instance, fitness data may be mined by an online provider to target advertisement to an individual based on readings collected by the individual's cell phone and stored by the provider.

Patients may feel that the data providers share with payers is more revealing of details than it needs to be. A common approach to addressing these problems is to subject information flows to the consent of the subject whose information is being shared. Consent is a cornerstone concept of medical privacy and provides a ready baseline for judging privacy protections in a given context. However, it does have important limitations.

First of all, rules for the protection of the public sometimes over-ride consent, such as laws for reporting gunshot wounds to law enforcement. There are also rules to demand or allow reporting personal data for purposes of medical research or public health. Second, patients are not always in a good position to judge whether a detailed instance of information sharing is too much or too little. Hence consent could decrease privacy by leading patients to make decisions to share data where they do not understand consequences that may be clear to persons with professional knowledge of the risks and benefits. In short, medical privacy and consent are deeply connected, but they are not equivalent.

Security threats in health care are an evolving concern. While in the financial sector the motives of an attacker are often clear, this is often not the case in health care. For instance, a phishing attacker wants to get types of private information that can be monetized in the online black markets, information like bank account passwords and credit card numbers. In the health care sector, where the data is usually not so clearly connected to a financial instrument, the motives are less clear.

Three factors need to be taken into account in predicting attacker motives and assessing security risks. First, health data often has associated administrative and financial data. For instance, patient demographics may well include enough information to get a credit card in the patient's name. Moreover, personal information can be used to file fraudulent claims. In countries where health care insurance is not available to everyone, there is an incentive for medical identity theft, in which an attacker gains the health care insurance benefits of another by impersonation. This obviously carries large risks for victims whose medical record is corrupted. Large-scale insurance fraud in the form of false billings is another common incentive.

Second, health data may become collateral damage in an attack. A computer virus is probably indifferent to whether it is infecting a home PC used for entertaining children versus a PC that runs a safety-critical process in a hospital. This threat is exacerbated by the regulatory review process for hospital equipment, which may slow the updating of software, hence preventing the rapid application of security patches. This could lead, paradoxically, to a situation where a home PC for children is more secure than the hospital PC. Detection of an intrusion in a hospital would result in the system being taken out of service until recovery is carried out.

Third, even if health data may motivate fewer attackers than in other sectors of the economy, it is often exceptionally critical to the safety of an individual and its corruption can be life-threatening. It may seem unthinkable that someone could deliberately consider corrupting health data, until it is done. We can look to instances like the 1982 poison Tylenol murders in Chicago as an unthinkable attack on the integrity of a health product to see that HIT systems are at risk to such extreme attackers. This sort of problem is likely to apply to IT contexts; for example, there was an incident in which the Epilepsy Foundation web site was used to upload images that cause seizures and migraines when viewed by some epileptics.

Access controls and audit

A key challenge for providers implementing HIT is to regulate access to patient information. The obvious step of establishing access controls to limit personnel access to a "by need" minimum is challenged by the complexity of clinical workflows and the high risk of denying access to key information, like drug allergies and so on, to personnel who might be involved in reacting to an emergency.

On the bright side, unlike rooms full of paper records, it is possible to trace, through electronic logs, which users look at which records so an auditor can use this information to detect abuses. There have been many examples of abuses that were caught in this way. Some involve access to the records of celebrities like athletes and actors; others involve incidents where, for instance, an employee of a provider accesses the record of a former spouse. These and other types of abuses are often addressed by investigations carried out after a complaint.

For instance, if an employee uses patient records to get credit card information, auditors can trace and identify the employee by carefully analyzing log data. This reactive procedure is increasingly inadequate because it does not scale up to new threats like large scale identity theft or to the increasing magnitude of the problem posed by the growth of providers and their connections through HIE. Research is needed to provide better automation so that large volumes of records can be examined by computer algorithms that are thorough and flexible enough to learn and infer threats quickly and feed experience from operational behavior back into preventative measures.

Two common strategies for addressing this problem provide steps in the right direction. One is to establish heuristics for common types of abuses, such as an employee inspecting the record of a patient with the same last name as themselves. A second is to set up rules which can be over-ridden in an emergency, a strategy sometimes called "break-the-glass" security. Heuristics suffer the problem that they cover only the types of abuses for which rules have been well-recognized and hence still have a reactive character. Break-the-glass has the problem that an overly restrictive set of rules may lead to so many instances of glass breaking that they cannot be meaningfully reviewed (Røstad and Øystein, 2007).

However, both strategies can be seen as contributing to a process that has developed more fully in other areas such as the financial services sector (credit card fraud detection) and messaging (spam detection). Adaptation to health care requires addressing issues specific to health care, such as the potentially high risks of a mistaken denial of access. The general idea is to develop systematic ways to learn quickly from experience and use this learning process to manage access rights in a kind of continuous quality control. Figure 9.3 illustrates the approach known as Experience Based Access Management (EBAM) (Gunter, *et. al.,* 2011), which can be applied to reconcile differences between an ideal access model and the enforced access control.

Access logs are used to measure differences between existing enforced controls and an ideal model for access rights. The ideal model represents the rules that should be applied. For many reasons these rules are only partially reflected in the enforced controls implemented by the electronic records system. However, information from the access log can be compared to the ideal model. This comparison can itself inform an engineered system, here called the expected model, which is used to learn and model legitimate accesses for purposes of improving enforced control and generating action items for organizational enforcement. Technologies that aid the development of an effective expected model have been accelerating in recent years (Boxwalla, *et. al.*, 2011; Chen, *et. al.*, 2012) and will soon be funding their way into practice.

Figure 9.3. Experienced-based access management (EBAM)

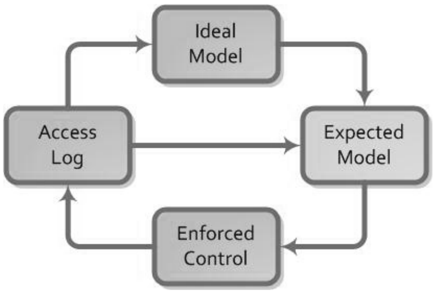

Source: Author.

Trusted base

Providers are struggling with rapid changes in the systems they need to secure. Early hospital computing systems used mainframe computers that could be accessed from terminals located in a hospital facility. This trusted base was relatively easy to secure until the Internet offered remote access, but standard enterprise protections such as firewalls and virtual private networks (VPNs) were accepted as being sufficiently effective. Now the situation is increasingly complicated by a range of technology changes.

Consider, for example, bring your own device (BYOD) arrangements in which employees put sensitive data on their own cell phones and tablets, the use of cloud services in which patient records are held by third parties, participation in HIE systems that move data between a changing collection of institutions, and the deployment of patient portals, which provide a new attack surface that can be assailed by unauthorized users for access to provider information systems. All of these changes redefine the nature of the trusted base.

Another area of concern is the rise of Advanced Persistent Threats (APTs), which entail sophisticated attacks, possibly supported by capable attackers like intelligence agencies. While there is currently no evidence that these attacks target health records, they are creating significant levels of collateral damage to EMR systems, especially when such systems are attached to certain types of targets like government and university networks. Such threats spur the need for greater attention to defining and maintaining the trusted base of health care systems.

Dealing with changes in trusted base requires careful risk analysis shows to determine which systems most need protection; proportionate measures can be taken for these systems. For example, a university hospital system that prepares records for certain types of research can de-identify records before they are shared with researchers; this provides risk mitigation in cases where the systems used by the university researchers operate at a lower security level than the trusted base of the hospital EMR.

Protection mechanisms established for the NIH-funded National Center for Biomedical Computing (NCBC) aiming to integrate Biology and the Bedside (i2b2) provides a case study (Murphy, *et. al.*, 2011). The system aims at a balance in which data that is subjected to more risk, such as data released to the public, is given proportionately more protection using techniques like de-identification. Encryption is a powerful tool for addressing challenges with trusted base. This is well-illustrated by secure transport protocols that allow data to be transmitted “in flight” over the Internet even when Internet routers are not trusted. This technique is used broadly for health care data, but less progress has been made on protecting data “at rest” in storage systems. Many examples of breaches of health care data have been of this kind. In particular, if the data stored on a laptop, removable media, or backup media is encrypted, then its physical theft is less of a loss. Similarly, if the data maintained by a compromised cloud service is encrypted, the threat of a privacy compromise is greatly reduced.

Research is needed to make such strategies efficient and convenient enough to enable their universal deployment, particularly to protect data at rest. General techniques can be applied to health records to achieve many goals, but there are also good ideas specific to health care (Benaloh et al., 2009).

Automated policy

A key challenge faced by many health care organizations (HCOs) is the need to share EHRs and exchange health data securely through HIEs such as those being set up by many states and regions in the United States, and through rapidly evolving partnerships with various business associates. Most HCOs must comply with a diverse set of policies, both internal and external, to exchange health data. The cost of ensuring compliance with these policies can sometimes be quite high due to the need for human policy experts and analysts to evaluate whether the organization is in compliance.

Current techniques to support health information exchange are too informal and manual to provide the desired efficiency and speed. For instance, if it is necessary to get an attorney to review and authorize each interstate data exchange by a provider in the United States, then a high level of exchange of EHR data will lead to a high level of expense (and delayed access). Enabling computers to settle policy decisions such as privacy compliance automatically can lead to reduced costs, improved care (though timely information exchange), and better support for secondary use of data.

Research is needed to determine reliable ways to formally express policies to enable fully automatic solutions. A benchmark that has been addressed by a number of studies (Breaux and Anton, 2008; DeYoung et al., 2010) is the formal specification of the U.S. Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule. We also require strategies to integrate and enforce formally expressed policies in common health care information architectures. Such advances will touch on other important areas like legal and medical ontologies and will inform the development of legal codes and consent management in the future.

Mobile health

A first concern for mHealth devices is how to secure the entire mHealth pipeline depicted in Figure 9.2. Some of the steps are familiar from current systems. For instance an enterprise laptop operating remotely will typically need to deal with a gateway device (like a wireless base station), the Internet, and an enterprise server. Typical solutions include security tunneling protocols like Transport Layer Security (TLS), IPsec virtual private networking, and wireless WiFi Protected Access (WPA). These techniques apply to mHealth devices as well, up to a point. For instance, a pedometer integrated into a cell phone can use whatever security is used by the cell phone to communicate its readings to a server. A more interesting problem arises when the pedometer is independent of the cell phone and needs a secure communication link of its own. Bluetooth seems like a logical choice for this sort of application and many of the first generation of mHealth sensors do indeed use a cell phone as a gateway device and Bluetooth security to protect the communications out of the mHealth device. This approach is challenged by problems related to secure pairing (because many medical devices will not have displays), privacy concerns about discovery mode, interference with other Bluetooth devices, and the scalability of Bluetooth to larger numbers of devices (Mare and Kotz, 2010). The growth of wearable computing devices has spawned the new subject of Body Area Networking (BAN) and security mechanisms appropriate to BAN will need to be developed.

Another area of concern is protecting the integrity and privacy of mHealth applications at nodes of the mHealth pipeline. For instance, a mobile application on a cell phone may share the platform with applications like video games downloaded from an application store. This type of sharing may expose significant safety risks particularly for actuator devices like insulin pumps. Yet another concern is the nature and motives of parties like EHR/PHR and analysts on the right side of the mHealth pipeline, many of whom will have business models that envision monetizing health data through data mining. For instance, a health device may offer a “premium plan” for people to share their data in a pool for comparison. This is not a bad thing by itself, but such intermediaries are likely to use personal data more freely than, for example, HIPAA entities are allowed to do in the U.S. and this could violate expectations for many mHealth clients.

One of the primary areas of current concern for mHealth is to identify requirements for privacy and security that are special to the space [Avancha *et. al.,* 2013]. Rules of the road for intermediaries need attention, but many of these can follow precedents like fair information practices. Other aspects seem newer. For example, there is an exceptional need to develop good isolation for applications on cell phones if mHealth applications are to run securely there. This problem is very similar to the trusted base issue of BYOD for cell phones in enterprise applications and it may be that the same security solutions can be used for both. However, there are instances where there is no clear analogy. For instance, the vulnerability of remote-controlled medical devices remains a concern since it has been shown (Halpern *et. al.,* 2008) that current wireless links are vulnerable to attacks on the integrity of widely used types of implants. One interesting direction is the use of “amulet” like auxiliary devices that provide security with good tradeoffs for these requirements (Gollakota *et. al.*, 2011; Sorber, *et. al.*, 2012) achieved, for example, by jamming unauthorized wireless communications with the medical device. In some cases there will be a desire to keep the existence of an mHealth device private. Potential examples include fetal heart monitors and devices associated with controlling addictions.

Identification and authentication

A long-standing problem in health care delivery is the risk of misidentifying a patient. Misidentifications cost lives, but procedures to reduce this risk are often cumbersome and may impede effective sharing of data between institutions. In addition to the problem of identification there is an emerging problem with authentication, that is, in proving identity. Inadequate authentication procedures are exploited by attacks like medical identity theft.

Increasing use of computer-based access diminishes traditional mechanisms of authentication like face-to-face meetings between individuals who know each other personally. This problem will become worse with the deployment of HIEs, which greatly increase the pool of people for whom identification and authentication are required within a single system.

While some of the problems in this area are non-technical policy concerns (like whether a national identity number system can be imposed) and many issues will be sufficiently addressed by broader Public Key Infrastructure (PKI), there is also a need for novel contributions. What is especially needed is a "science of identifcation and authentication" in which studies that involve the full gamut of regulatory, human factor, cryptographic, computer system, and other relevant considerations are subjected to analysis so that meaningful progress can be made and measured (Bonneau *et. al.*, 2011). Current research in this area needs to be expanded and integrated with operational approaches that can scale. For example, one of the earliest studies on PHRs in the U.S. used triple-factor authentication for both patients and health care professionals (Masys et al., 2002), an approach that is very secure but unlikely to be scalable for usability and maintenance reasons. By contrast, in the German Nationwide Health Information Technology Infrastructure (HTI) (Dehling and Sunyaev, 2012), medical professionals are given Health Professional Cards (HPCs) while a distinct class of Secure Mobile Cards (SMCs) are associated with institutions like hospitals and pharmacies. This division allows larger institutions to operate by delegation using SMCs so that the HTI authorities do not need to maintain authentication information for all employees.

Data segmentation and de-identification

It is widely recognized by both HCOs and government regulators that patients feel that some types of health data are especially sensitive. Examples include records related to mental health, drug abuse, genetics, sexually transmitted infections, and more. When health data is shared, there is a desire to transmit this information only when it is necessary. For example, a provider who needs immunization records may not need to see mental health notes. Interest in how to perform this kind of data segmentation has intensified with the growth of HCOs and the introduction of HIEs. However, there is little understanding of exactly how this type of segmentation can deliver meaningful privacy with acceptable impact on the safety and quality of care. Vendor products that claim to segment data may mislead patients and providers if they are poorly designed. De-identification can be viewed as a special instance of data segmentation in which information that personally identifies the patient is redacted or abstracted. The data segmentation problem needs some of the rigor that has been applied to the de-identification problem. In particular, we require ways to measure the tradeoffs between privacy, safety, and quality. These measures should be used to determine tradeoffs for specific segmentation technologies. For example, with de-identification there are measures of "diversity" that aim to quantify the level of privacy afforded by the identity-protecting transformations. There is a lively debate around the value of these measures and their practical application. By contrast we do not yet have any comparable measure that can be used to quantify the goal and effectiveness of removing, say, an item from a list of medications, as a protection against revealing a stigmatizing medical condition. It would be especially welcome to have some way to measure the impact that hiding information may have on care.

The de-identification problem itself also faces new challenges such as how to protect privacy of genomic data. Is genomic data like a lab result that can be treated like any other lab result or is it *intrinsically* identifying and therefore needs its own means of de-identification? New techniques are emerging in this area, for example, applications of cryptographic techniques that can be used to answer specific questions without revealing additional information. New research is required to determine information flows and privacy risks and to design sufficiently efficient protective measures.

Conclusions

The six privacy and security challenges described in this paper are not the only ones that face the area and they are overlapping in many instances. In particular, there are cross-cutting considerations that have not been listed explicitly. This concluding section will focus on two of these, namely the question of balancing benefits with costs and the impact of public policy and regulatory frameworks.

Balancing benefits with costs for security precautions is a long-standing challenge to justify expenditures for security protections.

In some instances there are clear quantifications that can be made. For example, at one time there were waves of virus attacks that had an impact on system integrity and availability; costs for these attacks could be calculated in terms of lost employee productivity and the need to assign IT staff to recovery and counter-measures. Security precautions are often developed with a "pierce and patch" philosophy where vulnerabilities are patched after they have been exploited by attackers.

One can reasonably expect to see this strategy being used for HIT just as it is in other areas, but the different circumstances and risks of HIT will often require special consideration. Implanted devices offer a good illustration. There is no evidence currently that there have been attacks on such devices, and basic security counter-measures like cryptographic authentication have costs. Since implants are limited by battery life and must be serviced surgically when the battery runs low, any cost of power must be strongly justified. Searching for the right balance is essential, but this balance is not something that can currently be plugged into a set of cost/benefit equations to get a solution. Aside from this general problem with security, the health and wellness space also faces challenges with balancing privacy costs and benefits. For instance, a patient may feel his privacy is protected by hiding a fact from a provider, but this hiding may lead to a misdiagnosis that wastes resources or harms the patient. On the other hand, if all information is routinely revealed then there are individuals who may decide to not seek care in a timely way from fear of disclosure, which may again lead to waste or harm.

In many countries health and wellness is deeply influenced by public policy and the incentives provided to private parties. Key stakeholders commonly include patients, payers (including tax payers), providers, vendors, and regulators. The kinds of research challenges that look important often depend on the interactions between these stakeholders. For instance, the particulars of research on how to model HIPAA make sense for the United States, which is governed by HIPAA, but not for the European Union, which has a different framework for privacy regulation that is not sector specific. However, there is common ground in some areas like the need for European vendors to comply to regulations for their US product sales, and the need for techniques to demonstrate compliance to regulations regardless of their local variations.

One thorny issue that will affect most health care regulatory systems is the scope of regulation for safety and security for devices. There are at least two driving issues. One is the growth in the types of devices that can satisfy a medical purpose previously dedicated to regulated medical devices. For instance, if a cell phone is being used as a stethoscope and stethoscopes are regulated, should cell phones also be regulated? Or should there be regulations only on the medical applications on cell phones that provide stethoscope functionality? Should this regulation also cover the relevant hardware on or associated with the cell phone that is involved in the stethoscope functions? Safety will undoubtedly be a leading consideration in these considerations, but security and privacy issues will have their own place. For instance, many of these medical capabilities will draw information into cloud services, and there is a question of their regulation. Also, the hosts for these types of information will begin to share it with regulated providers for the benefit of patients and perhaps for other reasons. How should this sharing be regulated?

Another interesting question that relates to cost/benefit assessment and regulation is the extent to which privacy and security should be considered an externality in the economic sense. That is, are privacy and security violations similar to pollution, where the true costs must be placed on responsible parties through regulatory controls? This is a common view for health care at providers where the economic incentives for privacy protections are often calculated in terms of fines to be avoided. Progress on cost/benefit analysis, regulatory incentives, and their combination will drive many aspects of privacy and security for health and wellness in the future.

References

Avancha, S., A. Baxi and D. Kotz (2013), "Privacy in mobile technology for personal health care", *ACM Computing Surveys*, 45.

Benaloh, J., M. Chase, E. Horvitz and K. Lauter (2009), "Patient controlled encryption: ensuring privacy of electronic medical records", in R. Sion and D. Song (eds.), *CCSW*, pp. 103-114, ACM.

Bonneau, J., C. Herley, P.C. van Oorschot and F. Stajano (2012), "The quest to replace passwords: A framework for comparative evaluation of web authentication schemes", in *IEEE Symposium on Security and Privacy*, pp. 553-567, IEEE Computer Society.

Boxwala, A.A., J. Kim, J.M. Grillo and L. Ohno-Machado (2011), "Using statistical and machine learning to help institutions detect suspicious access to electronic health records", *JAMIA*, 18(4):498-505.

Breaux, T.D. and A.I. Anton (2008), "Analyzing regulatory rules for privacy and security requirements", *IEEE Transactions on Software Engineering*, 34(1):5-20.

Chen, Y., S. Nyemba and B. Malin (2012), "Detecting anomalous insiders in collaborative information systems", *IEEE Transactions on Dependable and Secure Computing*, 9(3):332-344.

Dehling, T. and A. Sunyaev (2012), "Information security of patient-centred services utilising the German nationwide health information technology infrastructure", in *USENIX Workshop on Heatlh Security and Privacy*, Bellevue, WA, August.

DeYoung, H., D. Garg, L. Jia, D.K. Kaynar and A. Datta, "Experiences in the logical specification of the HIPAA and GLBA privacy laws", in E. Al-Shaer and K.B. Frikken (eds), *WPES*, pp. 73-82.

Gollakota, S., H. Hassanieh, B. Ransford, D. Katabi and K. Fu (2011), "They can hear your heartbeats: non-invasive security for implantable medical devices", in S. Keshav, J. Liebeherr, J.W. Byers, and Jeffrey C. Mogul, editors, *SIGCOMM*, pp. 2-13, ACM.

Grossmann, C.W., A. Goolsby, L. Olsen and J.M. McGinnis (2011), *Engineering a Learning Health care System: A Look at the Future: Workshop Summary*, The Learning Health System Series Roundtable on Value & Science-driven Health Care, The National Academies Press.

Gunter, C.A., D.M. Liebovitz and B. Malin (2011), "Experience-based access management: A life-cycle framework for identity and access management systems", *IEEE Security & Privacy Magazine*, 9(5), September/October, pp. 48-55.

Halperin, D., T.S. Heydt-Benjamin, B. Ransford, S.S. Clark, B. Defend, W. Morgan, K. Fu, T. Kohno and W.H. Maisel (2008), "Pacemakers and implantable cardiac defibrillators: Software radio attacks and zero-power defenses", in *IEEE Symposium on Security and Privacy*, pp. 129-142, IEEE Computer Society.

Mare, S. and D. Kotz (2010), "Is Bluetooth the right technology for mHealth?", position paper in *USENIX Workshop on Health Security (HealthSec*), August.

Masys, D.R., D. Baker, A. Butros and K. E. Cowles (2002), "Giving patients access to their medical records via the internet: The PCASSO experience", research paper, *JAMIA*, 9(2):181-191.

Murphy, Shawn N., Vivian Gainer, Michael Mendis, Susanne Churchill, and Isaac Kohane. Strategies for maintaining patient privacy in i2b2. *JAMIA*, 18:10-108, 2011.

Nissenbaum, H., (2011), "A Contextual Approach to Privacy Online", *Dædalus*, vol. 140, issue 4, September, American Academy of Arts & Sciences.

Røstad, L. and N. Øystein (2007), "Access Control and Integration of Health Care Systems: An Experience Report and Future Challenges", in *Proceedings of the 2nd International Conference on Availability, Reliability and Security (ARES 07),* IEEE CS Press, pp. 871–878.

Sorber, J., M. Shin, R.A. Peterson, C. Cornelius, S. Mare, A. Prasad, Z. Marois, E. Smithayer and D. Kotz (2012), "An amulet for trustworthy wearable m-health", in G. Borriello and R.K. Balan (eds.), *HotMobile*, p. 7, ACM.