Privacy is Healthy

Kelly Caine
Clemson University

Kelly Caine is an associate professor in the Human-Centered Computing division of the School of Computing and the director of the Humans and Technology Lab (www.hatlab.org) at Clemson University. Contact her at caine@clemson.edu.

Abstract: There are numerous privacy challenges specific to healthcare, ranging from patient expectations for confidentiality to sensors designed to collect health-related data that falls outside the bounds of traditional medical practice. All of these challenges make healthcare a unique environment when it comes to privacy. Learn what ubicomp researchers and practitioners can do to improve the state of privacy in a ubiquitous healthcare environment.

Keywords: pervasive computing, privacy, security, healthcare, mobile, Internet of things, data analysis

Health is inseparable from everyday life. Intuitively, we know that the environment in which we live, the stresses we experience, and our attitudes affect how healthy we feel. Increasingly, the research community is confirming that these factors—from our microbiome to our disposition—are just as important to living a full and healthy life as the genes we inherit. Everything we feel and experience has the potential to affect health. So, when researchers say that people need privacy for health information, what they're really saying is that people need privacy in everyday life. If health is inseparable from everyday life, then the privacy protections afforded to health information should necessarily extend beyond the clinical realm.

Privacy has always been a primary concern in ubiquitous computing. Mark Weiser stated in his original treatise introducing the world to ubiquitous computing, "hundreds of computers in every room, all capable of sensing people near them and linked by high-speed networks, have the potential to make totalitarianism up to now seem like sheerest anarchy."[1] He noted that privacy was key among the social issues yet unsolved in his vision for the future. Furthermore, he predicted that no purely computer science approach would succeed at solving the privacy problems introduced by ubiquitous computing; addressing social issues, including giving control to the individual, would be required. When we combine the ubiquity of health information in our everyday life with ubiquitous computing, the potential for privacy concerns become even more extreme.

Pre-publication version. Please cite as: Caine, K. (2016). Privacy Is Healthy. *IEEE Pervasive Computing*, *15*(4), 14-19.

Privacy Challenges Specific to Healthcare

There are numerous privacy challenges specific to healthcare, ranging from patient expectations for confidentiality to sensors designed to collect health-related data that falls outside the bounds of traditional medical practice. All of these challenges make healthcare a unique environment when it comes to privacy.

Patient Expectations

Two main concerns exist when it comes to patient expectations.

Patient-provider confidentiality. Patients believe that what they disclose to their healthcare provider is confidential and thus not revealed to any other parties. Of course, patients understand that there are exceptions, such as when a doctor gives a patient's prescription to a pharmacist, but the patient might not understand that, in turn, the pharmacist will likely know the underlying medical issue that calls for such a prescription. Furthermore, although patients know that everything they tell a provider is entered into an electronic health record (EHR), they might not fully realize how that information is shared with entities, ranging from the patients' own health insurance company to, in some cases, government officials, such as a state health department (see Figure 1).

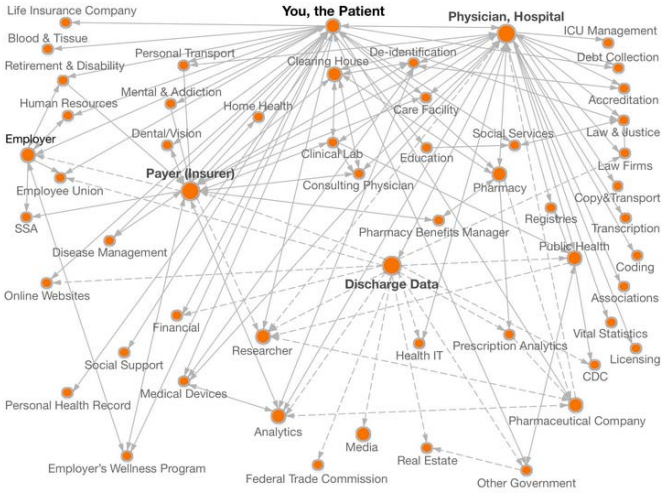

Figure 1. The health data map. It documents how personal health data can be shared with other entities. (Source: The Data Map, http://thedatamap.org/map2013/index.php)

Misunderstandings of HIPAA. The Health Insurance Portability and Accountability Act (HIPAA) is often misunderstood by patients, providers, and healthcare institutions (see the FAQ at www.worldprivacyforum.org/2013/09/hippaguudeindex). Patients often think that the privacy and security rules that are part of HIPAA protect all health information and apply to all entities. In reality, HIPAA only covers patient information kept by healthcare providers, insurance companies, and data clearinghouses and their business partners. It doesn't cover other entities that patients might reasonably include in their conception of "health information."

For example, HIPAA doesn't cover gyms, medical or fitness apps, occupational health clinics, fitness clubs, home testing laboratories, massage therapists, nutritional counselors, alternative medicine practitioners, or some urgent care facilities. Some entities, such as a grocery store, might be only partly covered, and although pharmacy information is protected, over-the-counter medication purchases from the store where a pharmacy is located aren't covered. Perhaps even more confusing is the fact that HIPAA specifically allows health information to be shared freely with entities from whom patients might want privacy protection. For example, although some patients think that HIPAA means a doctor is prohibited from sharing private, individual-level health information with an insurance provider or public health agency, the opposite is true; providers are specifically permitted to share a patients' health information with these entities.

In addition to only applying to specific entities, HIPAA also creates a false dichotomy between "protected" health information and "other" health information. In this way, it fails to protect lots of things patients consider health information. For example, data from fitness trackers, health apps, home-based health tests (such as paternity tests), do-it-yourself genetic tests (such as 23andMe; www.23andme.com), genealogy, or online portals where patients store their own health data surely seem like they should be covered by laws designed to protect the privacy of health information, but they're not covered under HIPAA.

Increasingly, in addition to using this information on their own, patients also want to share this "extra-clinical" information (data that originates outside a doctors' office, lab, or hospital) with healthcare providers. It's easy to see why patients would want to access and benefit from data collected during their daily life; extra-clinical information can provide a lot of information about a patient's current health. However, once this information is recorded, it can be used for purposes the patient never intended, or for purposes the patient specifically thought were covered by privacy protections.

For example, in one case, police used genealogy data from an online genealogy database to search for a match to DNA found at the scene of a crime.[2] Only after investigators tested an original sample of the man's DNA was he cleared of involvement of the crime; the data in the genealogy database had led investigators to an erroneous conclusion. Most likely, the people who used the genealogy service never intended to provide their private DNA information to police. Instead, they might have misunderstood or been unaware of the privacy policy and thought that their DNA information would be kept private by the genealogy service.

The Case for Privacy in Healthcare

Many of my colleagues have made the case for the need to improve the privacy protections afforded to health information. The case for privacy when consumer healthcare devices are pervasive is even stronger. Generally, there is a strong case to be made for keeping health information private. Notably, there are even cases in which a patient might want to keep some health information private from him or herself. The case for not knowing some health information is tricky and clearly individual.

A Diagnosis without a Test

In the same way that a placebo can make a patient feel better and experience better health outcomes, a diagnosis can, at least initially, make patients feel worse and potentially experience negative health outcomes. Of course, this is tricky, because, ideally, a diagnosis leads to treatment, which resolves or ameliorates the underlying cause of an illness.

Historically, the timeline for a diagnosis goes like this: a patient has a complaint, seeks medical attention, and has a test or series of tests performed. Then, a diagnosis is rendered and the patient is notified, usually in person by a healthcare provider (see Figure 2a). However, this is all changing.

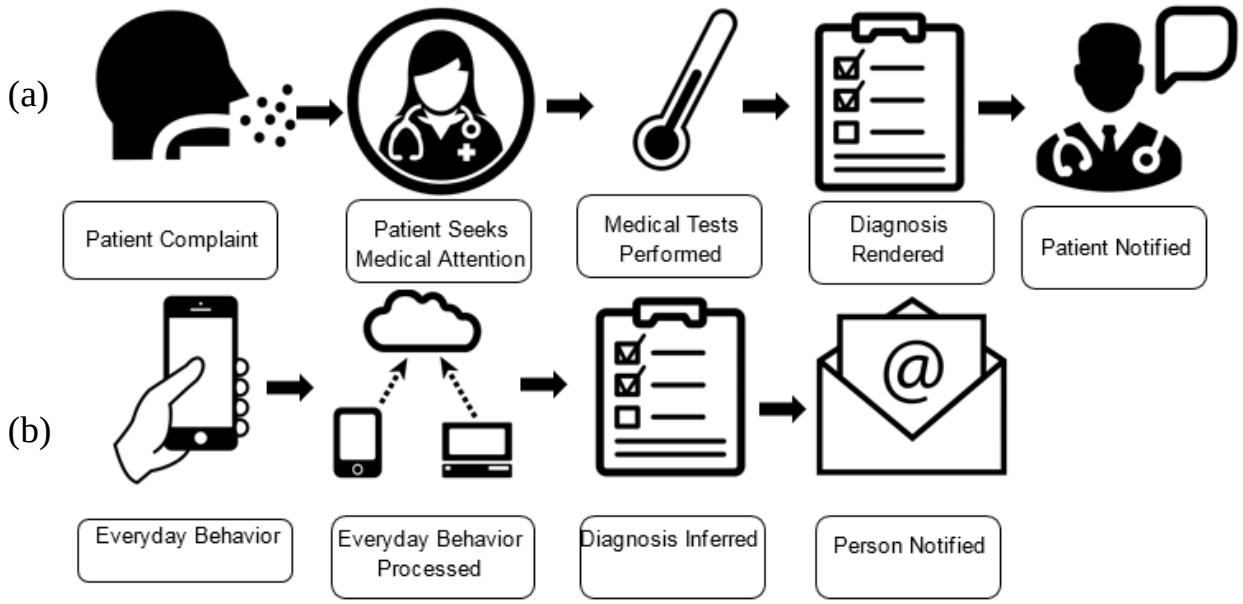

Figure 2. The shifting paradigm in the timeline for a diagnosis: (a) the historical timeline versus (b) the emerging timeline. Note that the patient initiative is much lower in the emerging paradigm.

Emerging techniques make it possible to suggest a diagnosis without the need for a traditional medical test; you don't even need to have a person to initiate interaction with a healthcare provider (see Figure 2b). Instead, in the emerging paradigm, a person engages in everyday behavior; then, that everyday behavior is sensed and processed, and a diagnosis is inferred from the data. How or whether a person is notified of the outcome has yet to be determined.

One well-known example of this emerging ability to sense health issues from everyday behaviors is Google's flu trends. Using individual user search input, Google was able to estimate flu outbreaks by geographic region.[3] Initially, these results were heralded as being as accurate as the Center for Disease Control's flu tracking, but up to two weeks faster. Although later results proved to be less accurate than originally thought,[4,5] this example illustrates the point that researchers can glean important population-level health information using individual-level everyday behaviors, without any traditional medical

tests required.

Using everyday behaviors to learn about individual-level health is also becoming possible. In the same way that Google flu trends was able to gain population-level information indicating flu outbreaks, Microsoft researchers have demonstrated the ability to identify cancer patients using individual online search patterns. Specifically, the team from Microsoft Research was able to identify up to 15 percent of patients who are thought to have later been diagnosed with pancreatic cancer based solely on search queries to Microsoft's search engine, Bing.[6] The researchers concluded that using individual behaviors, such as search queries, might be used as a signal that can predict a forthcoming pancreatic cancer diagnosis.

Other behaviors, in addition to what we specifically search for online, might also be harbingers of the onset of declining health. For example, although it's straightforward to consider how monitoring performance on online cognitive tests might be able to distinguish older adults experiencing cognitive declines from those on a more stable trajectory, what is less intuitive is that the mere cadence of typing could provide clues about cognitive decline. In fact, both test scores and typing cadence can predict cognitive decline.[7] Similarly, using gait-recognition technology via simple video recordings or an accelerometer in a phone or wearable, it's possible to observe changes in how a person walks and moves. Recognizing a gait change can be an indicator of the onset of dementia or other neurologic conditions.

Any information gathered about a person, as simple and ubiquitous as typing or walking, can potentially be used to draw health-related conclusions. Notably, these emerging diagnostic techniques don't require a patient to seek out medical attention or even to experience a symptom or seek the help of a healthcare provider. Rather, it's possible to diagnose people before they ever seek treatment. This is a giant shift in healthcare. The upside of being able to diagnose without traditional testing is obvious: we can catch diseases earlier—before a patient even experiences conscious symptoms—and provide earlier treatment, hopefully leading to better health outcomes. However, along with this advancement, and all the possibilities it offers for improving care and saving lives, comes huge questions and potential downsides that we need to consider, especially around patient privacy.

Seeking a Diagnosis vs. Learning You're Sick

Because a patient diagnosed using everyday behavior might not have experienced any complaints, medical professionals might more accurately be referred to as a "person" than a "patient". The person doesn't feel sick; hasn't decided to engage with the healthcare system; and is living life, potentially happy with the state of things. So when is it okay to interrupt this bliss?

What if the issue we can diagnose has no treatment? Do we still tell the person? What if the treatment carries a substantial risk? What if the treatment is very expensive or not covered by health insurance, so treatment for that person is unlikely? Can other potentially interested parties, such as life insurance companies or the persons' employer, access this information, regardless of whether the person knows there is a diagnosis? If a person didn't consent to being tested or notified of a potential issue, is it ethical to conduct the test in the first place? What if the test isn't really a "test" in the traditional sense of the word, but rather a picture that emerges about a person's health?

Furthermore, if we discover information about a person's health, how do we know when we should reveal this information? If a person can know sooner, should they always? These questions call for centuries of biomedical ethical principles to be reconsidered.

Many of these questions also all center on privacy. Privacy is a multifaceted concept including more than just a piece of data being shared or leaked without a person's consent. It also includes, as Thomas Cooley and then later Samuel Warren and Louis Brandeis put it, "the right to be let alone."[8] The psychological concept of privacy includes the right not to be intruded upon or bothered as we go about our daily lives. This is why, for example, people feel telemarketers, spam, and junk mail are invasions of privacy; they invade our daily lives without invitation. Shifting from a health system in which people specifically seek health information to one where they're diagnosed with no test will require significant consideration to avoid this sort of high-stakes privacy invasion.

The Dangers of Testing

We have some evidence from various cancer screenings that waiting to be tested might actually be better for health outcomes. Framed in terms of privacy, keeping private even from oneself and one's healthcare provider the potential diagnosis resulting from a test might be the best option. Recently, the guidelines for prostate and cervical cancer screenings were revised so that there are fewer tests or the tests start later in life. These changes are based on evidence suggesting that the tests detected cancers that the body might fight off naturally, the treatment for the cancer had severely negative side effects, or the cancer was unlikely to cause severe consequences (death) before the person died of other unrelated causes.

This should serve as a reminder that as we develop tests that use everyday behaviors to diagnose health conditions, we should be careful of over-testing and over-treating. The dangers for over-testing are different from over-treating. For over-testing, one of the major concerns is causing undue stress to individuals. This stress relates to the aspect of privacy in terms of not being interrupted or bothered in daily life; it doesn't relate to information being disclosed inappropriately or in an undesirable manner to other parties. How would having a constant health score made up of our daily activities affect our stress level? And how would that stress level cycle back to play out in our health outcomes? We need to answer these questions before we start diagnosing people who aren't seeking a diagnosis.

Another danger in testing—or rather in using ubiquitous health technology, especially wearables, to "incentivize" people to adopt healthier behaviors—is that the incentives are often ill-considered. Unfortunately, instead of incentivizing people to behave in a healthier manner, many of these programs increase the daily stress of participants, encourage them to cheat, or fail to achieve increases in positive health behaviors. An example is coercive policies from health insurance companies or employers offer "discounts" for filling out health questionnaire forms or for being tracked. Many of these programs aren't structured with true incentives but rather using coercive penalties. For example, some employer tracking systems essentially punish anyone who refuses to be tracked with a wearable health tracker by making those people who want privacy pay a fine or by not offering them a reduced premium.

The problem here is threefold. First, being constantly observed causes stress. We don't know the potential health consequences of the stress of being tracked by an employer

versus the potential for this tracking to increase healthy behaviors. Nor do we know whether there are better ways to incentivize healthy behaviors, such as offering a free gym membership instead of monitoring activities with a wearable health tracker.

Second, economic penalties, such as charging those who opt-out of being tracked for privacy reasons, might also cause stress or other health consequences to the individual. If people pay $50 more for health insurance each month because they opt-out of being tracked, they don't have those funds available for other health-related costs, such as healthier food, or other healthcare costs, such as a wellness check-up or preventative medication.

Finally, we know that some incentives, such as "step challenges," where employers give pedometers to employees and hold competitions to see who takes the most steps, unfortunately encourage some participants to cheat. Creative cheating strategies include putting fitness trackers on vibrating hand tools, a hamster wheel, or even a dog. These kind of ill-considered incentives induce people to "get the numbers" without actually leading to improved health.

Recommendations for Ubicomp Researchers

Although many of the questions I have posed here are perhaps suited for exploration by biomedical ethicists, others sit squarely in the ubicomp domain. There are things we, as ubicomp researchers and practitioners, can do to improve the state of privacy in a ubiquitous healthcare environment. And because we understand the totality of what can be sensed and inferred using ubiquitous technology, we should certainly do all we can to fight for privacy.

Build in Privacy by Default

My colleagues and I have been working to build privacy-enhanced health systems. One of the design principles we hold dear is to build in privacy by default. From the ground up, privacy is a design requirement. Privacy can manifest itself in various ways. In the Supporting Older Adults and their Caregivers Electronically (SOLACE) project,[9] my colleague Kay Connelly and I focused on ways older adults and their caregivers could co-manage the privacy of sensed activities of daily living.

For the Amulet project (https://amulet-project.org), working with colleagues from Clemson and Dartmouth, we've focused on building a secure mHealth wearable, where patients maintain control of the information that's collected and have a say about when and with which healthcare providers that information is shared. In these ways, we've built in privacy, so that users of these technologies stay in control of their health information, regardless of where it comes from.

Understand the Limitations of Privacy Protections

Understand that the technologies you are building might be gathering health data that is unprotected by current privacy laws such as HIPAA. A decent heuristic is that HIPAA (and some state laws) applies to clinical information but not information generated from commercial applications. Also understand that users probably don't understand this distinction; to them, all "health information" is protected.

Furthermore, understand that users don't read terms of service or privacy policies. If

you are assuming users understand these documents and rely on them to indicate that users agree that you can gather and share their health data, you are mistaken. Users aren't truly consenting to the use of their data via these instruments; rather, they're just clicking through without understanding the process or consequences.

Put Data in the Hands of Patients

A critical first step in building privacy into ubiquitous health systems by default is to put patient data in the hands of patients. From a privacy perspective, without easy access to their own data, patients have no way of knowing what exists and thus don't know who has access or who even has the ability to make choices about future access. Beyond privacy, patients should be able to easily access their own health data so they can understand their health situation (if they want to) and collaboratively make decisions about their future healthcare with their providers. We should design systems that assume that patients own their health information and have the right to first and exclusive access, if that is their choice. We should also work to find ways to let patients choose whose hands they want touching their health data, whether that data is HIPAA covered or not.

In addition to health data, patients should also have access to meta-information about how their health information is being accessed and used. One straightforward way to accomplish this is to offer a log of what health information is accessed by whom. Using visualizations combining temporal data, access logs, and health data types could help patients easily see any discrepancies between the access they expect providers to have and access they wish to forbid. In my own work,[10,11] and other excellent work studying the effects of providing patients with access to their health information,[12,13] generally patients prefer to be able to access their health information and control provider access.

Reaffirm the Patient-Provider Relationship

Another way ubicomp researchers can facilitate privacy in ubiquitous health systems is to reaffirm and recodify a privileged patient-provider relationship. Patients trust healthcare providers. They want to share their health information with the providers they trust.

Without realizing it, we (designers, computer scientists, insurance companies, government agencies, and lawyers) have let a system evolve in which patients can't reveal anything they want to keep "private" to their healthcare providers; this information is shared far beyond the bounds of historical patient-provider confidentiality by default. We should work to reaffirm and recodify a privileged patient-provider relationship. Patients should be able to reveal information to their healthcare provider in confidence, without the concern that it will be shared with anyone else. We can choose to design technologies in which privacy is the default, and patients must specifically choose when and with whom their information is shared.

Offer Healthier Options

Rather than monitoring what people do to punish those who behave in undesirable ways, focus on offering healthier options. How many of us want to be reprimanded by our employer when we haven't taken enough steps? On the other hand, how many of us would love to have the time and freedom in our work day to take a walk with a colleague or friend? Instead of being told our health insurance premium will be higher if we don't

drop a few pounds, why not offer a free gym membership, or free healthy meals at work?

Instead of creating technologies that track, monitor, surveil, and report undesirable behaviors, create technologies that encourage and enable healthier behaviors. But take care: even well-meaning programs that coercively incentivize “healthier” behaviors run the risk of being ethically questionable. Complicating matters here is that you must consider that what constitutes healthy behavior or healthy eating isn’t always straightforward. For example, although we once viewed fat as the enemy of a healthy diet, it might be that sugar is the real culprit. Similarly, a reduction in BMI or weight or increase in exercise doesn’t always equate to better health. Finally, realize that a desire for privacy in our everyday lives is a healthy behavior.

Privacy is healthy. From wanting a privileged and trusted relationship with a healthcare provider to maintaining control of when to seek health information about ourselves, it’s natural for people to want their health information private. The need to have space to contemplate our innermost thoughts and feelings without the judgment of others—that is, privacy—is integral to our health and well-being. Increasingly, health information can be collected ubiquitously, without our knowledge. Privacy for everyday information is key to maintaining health privacy in a world filled with ubiquitous technology.

Acknowledgments

I would like to thank Sheri Alpert, Jesus Favela, and Stephen Intille for their helpful comments on the manuscript. This work was supported by NSF grants 1513875, 1527421, 1314342, and 1117860. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

References

1. M. Weiser, “The Computer for the 21st Century,” *Scientific Am.*, vol. 265, no. 3, 1991, pp. 94–104.
2. J. Mustian, “New Orleans Filmmaker Cleared In Cold-Case Murder; False Positive Highlights Limitations of Familial DNA Searching,” *The New Orleans Advocate*, 12 Mar. 2015; www.theadvocate.com/new_orleans/news/article_1b3a3f96-d574-59e0-9c6a-c3c7c0d2f166.html.
3. S. Cook et al., “Assessing Google Flu Trends Performance in the United States during the 2009 Influenza Virus A (H1N1) Pandemic,” *PLOS One*, 19 Aug. 2011; http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0023610.
4. D. Butler, “When Google Got Flu Wrong: US Outbreak Foxes a Leading Web-Based Method for Tracking Seasonal Flu,” *Nature Int’l Weekly J. Science*, vol. 494, 2013, pp. 155–156.
5. D. Lazer et al., “The Parable of Google Flu: Traps in Big Data Analysis,” *Science*, vol. 343, no. 6176, 2014, pp. 1203–1204; http://science.sciencemag.org/content/343/6176/1203.
6. J. Paparrizos, R.W. White, and E. Horvitz, “Screening for Pancreatic Adenocarcinoma Using Signals From Web Search Logs: Feasibility Study and Results,” *J. Oncology Practice*, 2016; doi: 10.1200/JOP.2015.010504.

7. J. Kaye et al., “Unobtrusive Measurement of Daily Computer Use to Detect Mild Cognitive Impairment,” *Alzheimer's & Dementia*, vol. 10, no. 1, 2014, pp. 10–17.
8. S. Warren and L. Brandeis, “The Right to Privacy,” *Harvard Law Rev.*, 15 Dec. 1890.
9. I. Arreola et al., “From Checking on to Checking In: Designing for Low Socio-Economic Status Older Adults,” *Proc. SIGCHI Conf.Human Factors in Computing Systems* (CHI), 2014, pp. 1933–1936.
10. P. Swartz et al., “Patient Preferences in Controlling Access to Their Electronic Health Records: A Prospective Cohort Study in Primary Care,” *J. General Internal Medicine*, vol. 30, no. 1, 2015, pp. 25–30.
11. K. Caine et al., “Designing a Patient-Centered User Interface for Access Decisions about EHR Data: Implications from Patient Interviews,” *J. General Internal Medicine*, vol. 30, no. 1, 2015, pp. 7–16.
12. S.S. Woods et al., “Patient Experiences with Full Electronic Access to Health Records and Clinical Notes Through the My HealtheVet Personal Health Record Pilot: Qualitative Study,” *J. Medical Internet Research*, vol. 15, no. 3, 2013; doi: 10.2196/jmir.2356.
13. K.M. Nazi et al., “VA OpenNotes: Exploring the Experiences of Early Patient Adopters with Access to Clinical Notes,” *J. Am. Medical Informatics Assoc.*, vol. 22, no. 2, 2015, pp. 380–398.