Computerization, Discretion, Freedom

Sergey Bratus, Anna Shubina

December 31, 2015

Abstract

Surveillance of social networking, pervasive user tracking in hopes of reaping profits promised by "big data", and ubiquitous failure to secure stockpiled personal data went from being the concern of the few to making mainstream media. We've learned that what hurts privacy is also likely to hurt freedom. But, despite all these revelations, the worst and the most pervasive danger of computerizing our everyday lives has so far avoided public attention: that computers modify our behaviors related to discretion, professional autonomy, and, ultimately, moral choice.

Computerization changes every area of human activity it touches, by bringing new rules and new metrics. With enough of these at work, humans must act with an eye to not just what they do (or should do) in the actual real-world situations, but also to how it will look in the computer representation of it—and the latter are never complete. And when they disagree, one must either spend the extra time and effort "fighting the system", bend the rules—or give up.

In some areas such as the medical domain, "fighting the system" is a matter of life and death. Joe Bugajski credits a heroic ICU nurse for saving his life by fighting an electronic medical record system that, as he puts it, "nearly killed me". In a hospital emergency room nurses begin hoarding drugs to save patients' lives after a computerized drug dispenser had failed on them. Similar examples abound. Bending the rules is almost a given, and is not limited to life-critical domains. In fact, it has become so pervasive as to require systematic study.

Rules must be bent—but what if they can't be, or the risk to the employees is too high, and they have to give up?

Rules you cannot bend, when you clearly see you should, create helplessness. There is nothing more harmful to autonomy and freedom than helplessness, and computers are great (and guiltless) enforcers of rules; they are also great teachers of helplessness. Being forced to act against our best judgment at work is arguably much more damaging than knowing that we may be snooped upon in our playtime—at least for those of us who define what we are by what we do. Not only interfering with our daily professional activities cuts much deeper into our identity than an intrusively targeted ad, but it also arguably creates guilt and a feeling of helplessness beyond the Panopticon. Panopticon, after all, is obviously a prison, but a badly computerized workplace may be a version of the Stanford Prison Experiment, where one can be induced to play the warden quite unwittingly—because "the computer says no".

Whither Computerization?

The idea that widespread deployment of computer technology has the power to improve any given area of human endeavor appears to have taken deep roots in the public mind. Computer technology, once a concerted effort to embed it in an industry has been made, is assumed to be able to improve efficiency, lower costs, reduce waste, and so on.

Yet computerization changes every area of human activity it touches; the changes are both subtle and profound, as they modify human behaviors related to discretion, perceived competence, and, ultimately, moral choice. Indeed, introduction of computers must be watched very carefully by anyone concerned about freedom.

Computer Magic in the Public Mind

Proposals that rely on computing as a centerpiece of their promise do not tend to encounter much public scepticism—smart grid and the electronic health record infrastructure initiatives being the latest examples (compare these with proposals related to the good old "rocket science", which encounters much more sceptical reactions).

As it stands, the warnings of security professionals that our current commodity computer technology might not be ready for such massive endeavors are largely ignored and hardly ever reach the general public, much less receive its attention. Even cautionary stories of computer technology enabling dystopias ("Minority Report" and "Gattaca" among others) make the implicit assumption of enough "smarts" and reliability on the technology's part to sustain the prolonged or indefinite operation of the dystopia.

It appears that a message to the contrary does not resonate with the public mind; trustworthiness of complex computer technology—if not now, then in the nearest future—is assumed to be a trivial matter, a forgone conclusion. In the landscape of the public debate, sceptics have truly become "vox clamantis in deserto".

The Magic, Old and New

> "One of the chief duties of the mathematician in acting as an adviser to scientists ... is to discourage them from expecting too much of mathematics."
>
> *N. Wiener, "I Am a Mathematician, The Later Life of a Prodigy," MIT Press, 1964*

Norbert Wiener wrote this at the time when mathematical modeling and its applications were seen as magic that, when applied to almost any area of human activity by competent enough practitioners, would provide instant insights, improvements, or both.

This perception was formed by the powerful and obvious successes of operations research and similar applied disciplines. In the popular mind, it was also both part and justification of the vision of the scientifically organized society to come, in which wrongs would be equitably righted and burdens equitably distributed through science—and what better means to do so than mathematics? So great was this trust in mathematics, that even obviously unrealistic claims of unprecedented and unstoppable superiority of planned economies that claimed their planning to be based on mathematical methods were given wide credence, including by famous economists such as the Nobel laureate Paul Samuelson [1].

Wiener, the father of Cybernetics, was among those responsible for its broadly perceived promise. We would do well to heed his warning and apply it to the apparent and widely acclaimed magic of *our* days: computerization.

"Computers Can Improve Anything"

Belief in great powers of a formal device or method must first take root among enough engineers and others having to do with technology; only then it can be taken up by the general public. Computers have garnered this belief early on, so that already in 1979 DeMillo, Lipton, and Perlis repeated Wiener's warning with regard to computing.

More importantly, these authors offered a convincing explanation of the technologists' growing belief in the power of computerized computation:

> Since "symbols" can be written and moved about with negligible expenditure of energy, it is tempting to leap to the conclusion that *anything* is possible in the symbolic realm. This is the lesson of computability theory (viz., solvable problems vs. unsolvable problems), and also the lesson of complexity theory (viz., solvable problems vs. feasibly solvable problems): **physics does not suddenly break down at this level of human activity. It is no more feasible to construct *symbolic* structures without using energy than it is possible to construct *material* structures for free.**[1] But if symbols and material objects are to be identified in this

[1] Here and below the boldface emphasis is ours.

> way, then we should perhaps pay special attention to the way material artifacts are *engineered*, since we might expect that, in principle, the same limitations apply.
>
> —*Richard A. DeMillo, Richard J. Lipton, and Alan J. Perlis [2]*

It seems that the unrealistic expectations that Wiener warned about with regard to mathematical methods were revived by the promise of scaling up such symbolic computations with computers—as if the only limitations that kept the original magic from delivering were those of scale, not of natural laws.

In other words, using the magic of computers to get around the hard problems of the world is fraught with the same danger as calling on the magic of mathematical modeling had been in the 1960s: when limitations arising from natural laws get in the way, mathematics or computer science can, in time, *discover*, *distill*, and *describe* these limitations in their own formal ways, but not *overcome* them.

The hardness dictated by natural laws does not go away. It merely manifests itself in a different pattern of failures. In most applications of computing we tend to ascribe these failures to *exceptional circumstances* in which the programs fail; but these exceptions are also part of reality, and the program's failure to account for them is fundamental, not a matter of a quirk or bad luck. On the contrary, the supposed quirk is a manifestation of the natural complexity—which underlies the laws we had hoped to skirt for free with the use of computers.

It is tempting to attempt overcoming such failure by piling on more effort and more programming complexity. However, given the penetration of computers into all walks of life, this time the failures may have a bigger impact than just technical disappointments or even isolated technological disasters. These failures may fundamentally change human activities, as they are already changing some of the oldest human professions.

The Failures and Fallacies of Computerization

> "However, the assumption that the aforementioned benefits [safer care, reductions in administrative costs, improved clinical performance, better communication between patients and caregivers] are highly correlated with health IT has not been adequately tested and there are some indications that the features needed to acquire one benefit may actually frustrate efforts to achieve another. In particular, there is a growing concern that health IT designs that maximize the potential for administrative and economic benefit may be creating new paths to failure."
>
> *The New Institute of Medicine Report on Health IT and Patient Safety* [5],

It seems to me that the conclusions reached by this committee regarding computerization of health IT should have been a common sense *starting point* when an application of computer technology is being considered, for the following reason:

"Computerizing" a complex human activity means replacing human effort in some tasks by programs (in software or hardware), and altering other tasks so that they can feed into these programs.[2] The stated goal is usually better efficiency or precision. However, if the match between the replacement's logic, the other task changes necessary to accommodate it, and the actual work flow is not perfect, human time and effort will be spent to account for the differences—likely interfering with overall results.

Such replacement is particularly fraught with danger whenever desired performance of the task involves human discretion. There are many reasons for this. For example, the newly "computerized" task must work not just in its most common circumstances, but also in exceptional circumstances that require the performer to do something differently—the performer going off the script, applying discretion.

Discretion can be "shifted" to neighboring tasks as a part of their necessary adjustment to automation, but it cannot be simply eliminated by it. Without discretion, productive activity will break down at the first situations the program cannot handle—especially if the automated activity is sufficiently complex. The newly "computerized" part must work also in all those exceptional circumstances that create the system's complexity and require the human performer to do something differently.

[2]This applies to all human activities, from accounting to assembly by robotic arms: activities surrounding the automated task must also change.

Obviously, the nature and variation of the task must be well-understood and well-described to the programmers before the replacement of human performers is attempted. However, the amount of information required and the effort of gathering and organizing this information may be non-trivial and costly—potentially more costly than the actual improvement from computerization.

Replacement can have unexpected effects on other related tasks, and especially on the human involved in the computerized task when human involvement is still required. Should the program logic fail to exactly match what needs to be done, the human will find himself working around or *"fighting the system"*, spending effort and, possibly, compromising his performance on other tasks.

In particular, the human will have to mentally compute the difference between what he would normally do and what the program will do for him and to find a way to get that difference done in the new environment, which does not account for the necessity to get it done.

Fighting the system is especially hard in the naturally busy and distracting work environments where following a standard process is relied upon for reducing errors. The temptation to follow the computer's scripted idea of this process whether or not it is warranted in the particular exceptional circumstances may be strong; breaking off the script may require considerable mental effort at the time when it may be the hardest to muster. In other words, in order to bend the rules, one must first notice the need to do so—and it may be hard to notice while being driven by a computerized script.

Bending Rules and Scripts to Save Lives

In computerized medicine, "fighting the system" can be a matter of life and death. Although doctors and hospitals are reluctant to share stories of computer failure—and, in many cases, are barred from it by contract [3]—case studies of computer failure in the medical domain are beginning to build up.[3] We will look at just two such examples, which nevertheless perfectly illustrate the perils of computerization.

The first case study [7], documented by emergency room doctors, describes a failure of an automated dispensing unit for drugs in a hospital's emergency room:

> A critically ill patient presented to a busy emergency department (ED) [..] Intravenous access was obtained and a variety of pharmacologic agents were ordered. The resuscitation nurse went to obtain medications from an automated dispensing unit (ADU), part of a computer-based dispensing system in use throughout the hospital. He found an uninformative error message on the computer screen ("Printer not available") and an unresponsive keyboard. The system did not respond to any commands and would not dispense the required medications.

It was later determined that the dispenser's unresponsiveness was due to it being overwhelmed by a "storm" of network messages between the central computer managing the ADU units and the units themselves—essentially, a denial-of-service condition. This condition was created by a failed software upgrade. Undertaken two weeks before the incident, this upgrade of what was considered a reliable, mature system led to a complete loss of function. It became necessary to repeatedly enable and disable the "safety interlock" feature that affected the ADUs—and this is what apparently created the "message storm" on the day of the incident.

Naturally, the hospital had a back-up plan in case of the ADU mechanical failure: drugs necessary for management of cardiac arrest were available in a sealed chest, a "crash cart", for just such a case. However, it did not contain the drugs required to *prevent* the patient from going into cardiac arrest:

> It did not occur to the planners that cases such as this one—not (yet) in cardiac arrest, but with highly time critical need for drugs—might occur. No scenario-based planning was done, which might have generated example cases that could have led to anticipatory changes in the crash cart (for example, stocking additional drugs that might forestall cardiac arrest). The organisation at this time was in a severe financial crisis, and the organisational leadership seemed blinded by the potential for savings represented by the ADU system. Objections on safety grounds tended to come from nurses or emergency physicians, who were not part of the formal planning team, and

[3] Medical malpractice court cases are another source coming to light; the recently published overview [4] is fascinating—and scary.

were tagged as obstructionist, non-team players, so their objections were treated as theoretical at best and specious or manipulative at worst.

Luckily for the patient, some of the required drugs were found on top of another ADU unit, as the ED staff, alerted by the nurse, went looking for them. More were phoned for and brought by a runner from the hospital's pharmacy. The patient received the drugs, and her condition improved.

So even though computer failure was taken into account, and exceptional circumstances were seemingly provided for, these circumstances have not been fully anticipated and covered. Those who understood it best, adapted—by bending the rules:

> The organisational response to the event is telling. Parts of the organisation believed the incident showed that the system was safe, since the nursing and pharmacy staff were able to overcome the problem and since no harm resulted. **Nurses, on the other hand, began hoarding drugs as they did not trust the system**, thus subverting one of its major organisational goals (inventory control).

The second case was documented by a patient, Joe Bugajski, a systems integration and data interoperability specialist, in a story strikingly called *"The data model that nearly killed me."* His ordeal began with a trip to the Emergency Room (ER) of a hospital sharing an electronic health record system with that of his allergist's office—a choice prompted by this sharing, the benefits of which utterly failed to materialize.

> During the last week of January 2009 a faulty electronic, networked, health information data model nearly killed me despite its vaunted status as a component of two state-of-the-art, health information systems at two of the world's most advanced medical facilities.

Although he came to the ER and then to the Intensive Care Unit (ICU) with what he believed was a detailed medical record containing the crucial information about his condition, this information, as it turned out, was never available to his doctors or nurses at the right time. Instead, most of the medical professionals who interacted with him were driven by the scripted processes that had them ask questions, take measurements, enter them into the software—ad nauseam—without actually addressing his condition:

> Throughout my stay, I was hooked to network attached monitors that incessantly sounded alarms to which no one responded. I was asked 11 times to repeat my medical history, medication, and allergies to as many different medical professionals. I was seen by seven doctors each of whom asked me similar questions. Five doctors were never to be seen again. All doctors mumbled something about putting their findings into the hospital's electronic records system—most did not according to ICU nurses. No one read my allergist's detailed report about my condition and health history. As I moved from ER, to an ER holding room for admitted patients, back to ER, and to and fro other departments for tests, and finally to ICU, I was visited by nurses and technicians who pushed laptops mounted on wheeled sticks. They checked my vitals; asked me questions about my history, medications, and allergies; and entered findings into the hospital's electronic medical record using the laptops mounted on wheeled sticks.

This description sounds surreal. How could it be that the single crucial piece of information for saving the patient was being continually missed? It appears that it was so because it was lost in the stream of other information that the computerized system demanded but could not prioritize—whereas a human professional free of the system's distractions easily would.

Most of these medical professionals obviously never realized that what they were doing was not in the patient's best interest. Without the computer system to drive their actions and interactions, they likely would have focused on the patient's description of the issue that sent him to the ER in the first place—but their work flow has been heavily modified by the software, and they apparently failed to conceptualize how much it has changed, or could not do anything about the changes. They went along with the system rather than fighting it—perhaps were not even aware of the need to fight it.

It took a person determined to fight the system to end the phantasmagoria:

> One heroic medical professional, the first nurse I met in ICU, worked to create a consistent record of my condition, allergies, and medications in the hospital's electronic health information system. She spent over one hour searching for previously entered data, correcting errors, and moving or reentering data. She argued with one doctor whose concurrent access to the hospital's system blocked my nurse's access to my information. She called the hospital's pharmacy repeatedly to get my medications delivered. She met and called doctors several times. She even convinced one doctor and a pharmacist to come to my room to resolve data errors in person. Despite these heroic efforts, I never received correct medications during my stay. Indeed, my wife snuck one of my inhalers into my room. After I used it, I finally began to recover.
>
> At one point during my battle with illness and electronic healthcare data, the only asthma medication that had kept me alive began to wear off. I knew that if I did not receive the right dose within an hour or so, my condition would deteriorate rapidly and I would die. This critical information I had repeated 9 times to doctors and nurses who recorded it in my electronic health record. They promised that I would receive the medicine when it was time. That time came and went. My lungs began to scream with pain. My respiration rate accelerated. My breathing became more labored. I was crashing. I begged the doctor who next stopped to check my condition for her help. She said she would authorize the prescription. The heroic ICU nurse stopped by my room, checked my electronic records, but she could not find the prescription. She then ran to find a doctor to authorize my medicine. She succeeded. I received the medicine. I lived.

Notice how in Joe Bugajski's case it was following the computerized process that put him in danger, breaking the rules that alleviated his condition, and fighting the process that saved him.

We believe it is such experiences that led a doctor we chanced to talk to about her experiences with electronic systems to tell us that the day she and her colleagues practiced the best bed-side medicine—doctors talking to patients and nurses, nurses responding to patients, etc.—was the day their IT center had a fire and was offline. The rest of the time, she added, we do the best we can.

A World of Human Circumvention

Bending the rules is almost a given in a computerized workplace—and is not limited to life-critical domains. Getting around the nominal rules as enforced by a computer to "simply do one's work" is becoming an inherent part and shared lore of many professions—or even a core competency.

In fact, it has become so pervasive as to require systematic study, a *science of Human Circumvention* [10]. As organizations find that their well-intentioned users continually circumvent security controls—established to achieve some organizational goals—for no other reason than furthering other and often more important goals of the same organization, a general way of thinking about such practices is needed. It must include both conceptual models for why they happen and metrics of how they happen.

A promising approach for understanding circumvention is characterizing its root cause: the mismatch between the security control designer's idea of the processes it applies to and the reality as seen by workers in the trenches. Often, "what really happens in the trenches doesn't match the technology's underlying assumptions or even purposes [8]. The designer's model that maps the domain to technologies and controls fails to preserve important structure of the domain; the map is a *mismorphism*. Mismorphisms may "lie at the heart of circumvention, because they characterize the scenarios that frustrate users—and often the resulting circumvention itself." [9]

Circumvention is often a symptom of the professional discretion necessary for the production process being unwittingly damaged by a control. I am indebted to Eleanor Saitta for the following observation: *Many organizations tend to view discretion as a bug, rather than an inherent part of their production process.*

Why do organizations underestimate the amount of discretion involved in their activities ("computations")? This may be an organizational counterpart to overestimating one's competence: organizations underestimate the complexity of what they do and overestimate their rules and procedures.

In order to analyze shifts of discretion, organizations must first understand how much they depend on discretion and what role it plays in their overall computation. When they lack such understanding, they tend to overestimate the returns of computerization. A rise in circumvention may then serve them as an early warning of the vaunted computer systems actually damaging their production process.

The Perfect Bureaucrat

The inability of human-replacing computers (or "robots") to make proper exceptions has long been a theme of science fiction (and likely helped create the phenomenon of science fiction as such). The inspiration behind these dystopian visions, however, was not merely fear of the unknown or of the technological future sweeping in too fast. It was a natural extension of the mundane experience of a rule-driven computation in which human discretion is limited by circumstances or by design: the bureaucratic process.

The computer is nothing but a fast, implacable, *zero-discretion* bureaucrat. In fact, this is precisely how Richard Feynman chose to describe it in his "Lectures on Computation". Any discretion and context awareness in the system can only come from human overrides.

As computer security professionals—that is, specialists in computer-related *failures*—it is our duty to explain the drawbacks of computerized decision making. Perhaps we could do it as follows.

- The typical user focuses on "computers" (users imagine shiny PCs, tablets, and other cool and impressive devices), whereas the real issue is "computation"—what the computer is supposed to do. The computer can only execute instructions, and these instructions can only come from humans—and must start as instructions given to other humans before they become computer code. After all, the programmers must be told what to program.

- Once we focus on computation rather than "computer", we realize that it predates the computer as such—as Ross Andersson argues, massive computations in the sense of rule-guided, multi-step procedures, with almost no discretion on the part of low-ranking participants were actually being performed by roomfuls of clerks, bureaucratic hierarchies (in China, for thousands of years), and even armies.

- A significant part of the public is familiar with bureaucratic procedure and understands the drawbacks created by strictly following rules regardless of context and zero-discretion policies. The public should be encouraged apply this understanding to computerized environments, especially where humans who could listen to a reasonable explanation and make an adjustment are too busy and hard to reach.

Discretion and Freedom

Rules you cannot bend, when you clearly see you should, create helplessness. There is nothing more harmful to freedom than helplessness. But computers are great (and guiltless) enforcers of rules, and therefore great tools to teach helplessness. Therefore they must be watched very carefully by anyone concerned about freedom.

Behind a great many ideological rifts is the difference of belief in the degree of other humans' competence, and, therefore, in the value of discretion. The same is true for managerial styles.

One recurrent attitude is that all humans (possibly excluding the thinker) are better off with a stricter set of rules to define their daily activities. The damage that unbendable, zero-discretion rules clearly do in some cases is believed to be compensated—on average—by the good they create across the society, which otherwise would suffer more damage from poor application of discretion and lack of control. In short, the underlying conviction is that humans, on average, cannot be trusted with discretion, for their own good and efficiency, and the more of their material activities could be described as strict procedures to follow, the better.

From this position, introduction of additional (expert-designed) rules and controls to see that these are obeyed is by default seen as beneficial. It is seen as safer to err on the side of adopting technologies that introduce new possibilities for rule-making and enforcement: book-keeping, auditing, bureaucratic procedure, registration, certification. Computerization is merely a recent addition to this list.

However, for the society to fully take advantage of these technological improvements, humans must also learn (or be trained, or obligated) to delegate their decisions to certified procedures and professionals trained to apply the rules. Depending on the rule system and the activity, exclusive delegation may even be *required* for the stability of the rule-based mechanism, or at least for the perceived fairness of its outcomes.

The best kind of delegation, is, of course, voluntary, as it does not require dealing with the negative consequences of obvious coercion. The feeling of helplessness is a strong motivator for voluntarily putting faith into special people and magic procedures (and even for hostility towards skeptics). Thus, helplessness is very handy for those whose modus operandi for improving the world is to reduce human discretion.

Challenges the Future Holds

The accumulation of rules in rule-based artificial intelligence systems is known to practitioners to make adjusting or retraining them barely tractable. Often, new rules are introduced to compensate for effects of previously introduced ones or their combinations; the overall effect is more complex interactions and reduced tractability. Similar effects plague organizational policies: rules and checks that may have served to compensate for a past mistake or simply to make someone's job easier pile up over time—and may outlive the circumstances that called them into being. Accordingly, organizations accumulate informal lore on rules that may be bent and forms that may be skipped.

However, computerization creates environments where rules can no longer be informally bent or bypassed at the worker's discretion; what's written in code may as well be written in stone. Worse, the apparent ease of manipulating computerized rules may encourage more eager rule-making—but in complex systems such ease is a dangerous illusion, due to the above mentioned tractability problems.

Indeed, computerized environments are seeing a different culture than the lore of allowed discretion—that of "fighting the system", or "tricking the system" in the old bureaucracies, notorious for their Kafkaesque intractability.

This presents a new and not yet well-understood challenge to designers of systems intended to "computerize" human activities to previously impossible degrees and levels of detail: they must understand the role of discretion in these activities and make sure to accommodate it. Such accommodation will certainly require new methodology for understanding both processes and the effects of their uniformization, intentional or incidental to computerization.

It is debatable to which extent a system's designers are liable for its misuse due to misunderstanding of its inherent limitations and overenthusiastic misapplication. Yet it is clear that such misuse will eventually reflect on the designers and technologies, and therefore designers must think, e.g., of how to avoid empowering compulsive and excessive rule-making—just as they currently think of preventing other kinds of user error. This, in turn, will require studies of systematic organizational errors due to the advanced technology-induced illusion of control and overestimation of the rule-makers' competence.

The hardest of all will be the challenge of predicting and offsetting broader societal impacts of ubiquitous limitations of discretion. We hardly have a clear idea of the limits beyond which shrinking of individuals' discretion becomes a threat to personal autonomy, turning individuals into submissive and uncritical rule-followers; however, we know that technically enforced rule-following in well-organized societies can be the most slippery of slopes.

The view of the computerized future as a kind of a Panopticon, where good deeds are not done, and fellow humans are not helped because of how the deed may appear to the ubiquitous overseers, is bleak enough. Yet another, even more insidious danger lurks.

Panopticon, being a shared prison, cannot completely eliminate the empathy between its prisoners, nor the shared spirit in which the overseer could be challenged by them. However, another famous experiment has been remarkably successful at eliminating such empathy far beyond the expectations of the experimenters and the conscious self-images of almost everyone involved: the Stanford Prison Experiment.

We may believe ourselves to be forewarned against playing the wardens in that kind of a prison, as we may consider ourselves forewarned about obeying evil authority, "just following orders". Yet computerization that chips away at our discretion is perfectly positioned to place us in the warden roles—by "just following rules"—more easily than we realize, and with less guilt in every particular instance.

We must, therefore, view those aspects of computer technologies that limit human discretion with utmost

caution—and, amid enthusiasm for their obvious benefits, it is our primary professional duty to also warn the society of their less obvious dangers to personal autonomy and freedom.

References

[1] *The Poverty of Samuelson's Economics*, Alan Ebenstein, http://freedomkeys.com/samuelson.htm

[2] *Social Processes and Proofs of Theorems and Programs*, Richard A. DeMillo, Richard J. Lipton, and Alan J. Perlis, Yale University Technical Report TR-82, p.7, http://www.cs.yale.edu/publications/techreports/tr82.pdf

[3] Darius Tahis, Politico, *Doctors barred from discussing safety glitches in U.S.-funded software*, 09/11/2015, http://www.politico.com/story/2015/09/doctors-barred-from-discussing-safety-glitches-in-us-funded-software-213553

[4] M.L. Graber, D. Siegal, H. Riah, D. Johnston, K. Kenyon, *Electronic Health Record-Related Events in Medical Malpractice Claims*, Journal of Patient Safety, Nov. 2015, http://www.ncbi.nlm.nih.gov/pubmed/26558652

[5] *The New Institute of Medicine Report on Health IT and Patient Safety*, http://spectrum.ieee.org/riskfactor/computing/it/new-institute-of-medicine-report-on-health-it-and-patient-safety

[6] J. Bugajski, *The Data Model That Nearly Killed Me*, http://securehealth.freshdefense.net/content/data-model-killme.pdf

[7] R.L. Wears, R.I. Cook,*Automation, Interaction, Complexity, and Failure: A Case Study*, http://www.dcs.gla.ac.uk/~johnson/complexity/Proceedings/Wears.PDF

[8] J. Blythe, R. Koppel, S.W. Smith, *Circumvention of Security: Good Users Do Bad Things*, http://www.cs.dartmouth.edu/~sws/pubs/bks13.pdf

[9] S.W. Smith, R. Koppel, J. Blythe, V. Kothari, *Mismorphism: a Semiotic Model of Computer Security Circumvention*, http://www.cs.dartmouth.edu/~sws/pubs/skbk15.pdf

[10] T. Xie, J. Blythe, R. Koppel, S.W. Smith, *Science of Human Circumvention of Security*, http://publish.illinois.edu/science-of-security-lablet/science-of-human-circumvention-of-security/