

Mismorphism: a Semiotic Model of Computer Security Circumvention (Extended Version)

Computer Science Technical Report TR2015-768

Dartmouth College

S.W. Smith

Dartmouth College
`sws@cs.dartmouth.edu`

R. Koppel

University of Pennsylvania
`rkoppel@sas.upenn.edu`

J. Blythe

University of Southern California
`blythe@isi.edu`

V. Kothari

Dartmouth College
`Vijay.H.Kothari.GR@dartmouth.edu`

March 2015

Abstract

In real world domains, from healthcare to power to finance, we deploy computer systems intended to streamline and improve the activities of human agents in the corresponding non-cyber worlds. However, talking to actual users (instead of just computer security experts) reveals endemic circumvention of the computer-embedded rules. Good-intentioned users, trying to get their jobs done, systematically work around security and other controls embedded in their IT systems.

This paper reports on our work compiling a large corpus of such incidents and developing a model based on *semiotic triads* to examine security circumvention. This model suggests that *mismorphisms*—mappings that *fail* to preserve structure—lie at the heart of circumvention scenarios; differential perceptions and needs explain users’ actions. We support this claim with empirical data from the corpus.

1 Introduction

Users systematically work around security controls. We can pretend this doesn’t happen, but it does. In our research, we address this problem via observation and grounded theory (Bernard and Ryan, 2010; Charmaz, 2003; Pettigrew, 2000). Rather than assuming that users behave perfectly or that only bad users do bad things, we instead observe and record what really goes on compared to the various expectations. Then, after reviewing data, we develop structure and models, and bring in additional data to support, reject and refine these models.

Over the last several years, via interviews, observations, surveys, and literature searches, we have explored the often tenuous relationship among computer rules, users’ needs, and designers’ goals of computer systems. We have collected and analyzed a corpus of hundreds of circumvention and unusability scenarios. We categorized 296 examples of these “misunderstandings” and the circumventions users undertook to accomplish their needed tasks. We derived the examples from 285 different sources and categorized them into 60 fine-grained codes. Because several examples reflect multiple codes, there were 646 applications of the codes linked to the examples; e.g., the example of a woman with a hysterectomy listed in the current record as having an intact, normal womb was coded as: 1. A copy-and-paste issue (because the “current” record reflected an earlier examination from before her recent surgery); and 2. The IT not representing the reality. Most examples had only one or two codes associated with them; some had as many as four.

Semiotic triads, proposed almost a century ago (e.g., Ogden and Richards, 1927), offer models to help understand why human agents so often circumvent computer-embedded rules. The triads reflect the differences and similarities among: a) what the speaker/listener is thinking, b) what words or symbols are used to convey those thoughts, and c) what is the reality or the thing to which they are referring. We suggest that these triads provide a framework to illuminate, organize, and analyze circumvention problems.

This Paper In this paper, we present these ideas and support them with examples from our corpus. Examples where we don’t cite a source came from interviews with parties who wish to remain anonymous. As we are working on developing a typology rather than supporting a hypothesis, many of the usual factors in confirmation bias to do not apply. Section 2 presents our model of how users’ actions will often differ from the expectations of the security designers. Section 3 hypothesizes how our model of differential perceptions and needs explains users’ actions and non-linear/non-monotonic responses to increases in turning the security knob higher. Section 4 and Section 5 then supports these hypotheses with data item from our corpus. Section 6 considers some related work, and Section 7 concludes.

2 A Semiotic Model for IT Usability Trouble

In a previous paper (Smith and Koppel, 2014) that organized an earlier corpus of usability problems in health IT into a coherent typology, we considered three sets:

- the mental model of the clinician working with the patient and the health IT system;
- the representation of medical reality in the health IT system;
- and the actual medical reality of patients;

Usability problems organized nicely according to mismatches between the expressiveness of the representation “language” and the details of reality— between how a clinician’s mental model works with the representations and reality.

Somewhat to our chagrin, we discovered we were scooped by almost a century. In their seminal 1920s work on the meaning of language, Ogden and Richards (1927) constructed what is sometimes called the *semiotic triad*. The vertices are the three principal objects:

- What the speaker (or listener/reader) *thinks*
- The *symbol* they use
- The actual item to which they are *referring*

Much of Ogden and Richard’s analysis stems from the observation that there is not a direct connection from symbol to referent. Rather, when speaking or writing, the referent maps into the mental model of the speaker and then into the symbol; when reading (or listening), the symbol maps into the reader’s (listener’s) mental model, which then projects to a referent, but not necessarily the same one. For example, Alice may think of “Mexico” when she writes “this country,” but when Bob reads those works, he thinks of “Canada”—and (besides not being Mexico) his imagined Canada may differ substantially from the real one.

As we now consider a new corpus of scenarios in security circumvention and other authentication misadventures, we hypothesize that this framework will also apply. We have a set of IT systems. Each system serves a set of users, and mediates access between these users and a cross-product of actions and resources. Each system has an IT administrator who worries about the security configuration—as well as users who worry about trying to use the resulting system for their actual work. For different systems, the user sets are not necessarily disjoint.

The interaction between the reality, the IT representation, and the mental models correspond to the vertices in Ogden and Richards’ triad:

- *Thought*. The *mental model* a party has about the actions users can and cannot (or should and should not) do with resources.
- *Symbol* (i.e. *configuration*): The representation of security policy within the IT system itself; the built-in functionality of the IT system, intended to express the correct workflow. (Here, we mean policy as the actual machine-actionable expression of admin intention, not a published instructional document.)
- *Referent* (i.e. *reality*) The actions users can and cannot do with the resources, in reality; the de facto allowed workflow.

Figure 1-a sketches this basic triad. In this framework, the primary mappings are counterclockwise:

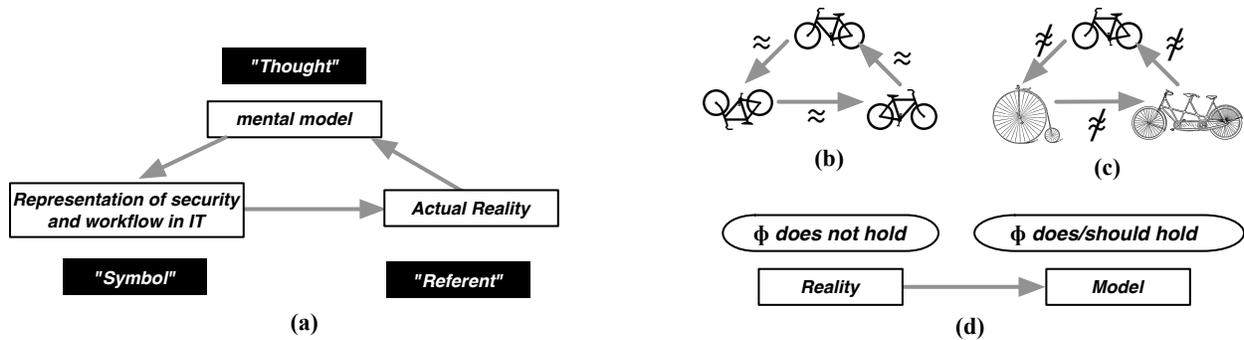


Figure 1: (a) The basic Ogden-Richards triad, moved into 21st-century IT; the arrows indicate the main direction of mappings. (b) Standard semiotics considers structure-preserving mappings between the nodes of the triad; (c) in circumvention semiotics, we think about mappings that fail to preserve structure. (d) E.g., in a standard mismorphism scenario, the generated reality fails to embody a property the user regards as critical.

- *Referent* \rightarrow *thought*: the admin constructs a mental model of what she imagines is the actual enterprise workflow requirements.
- *Thought* \rightarrow *symbol*: the admin reasons about security and work goals and construct a system configuration that she believes achieves these goals.
- *Symbol* \rightarrow *referent*: this configuration in practice then generates some actual reality.

Thanks to the connection of IT and reality, we now have a direct symbol-referent connection, improving on (or at least complicating) the merely linguistic world Ogden and Richards explored. Note however, that ordinary users also participate in this triad, and that mappings in the other direction can also be interesting: e.g., investment bankers trying to infer which of their entitlements are actually necessary in their daily job (symbol-thought, then thought-referent).

3 Extending this Model to Security Circumvention

To illustrate the role of semiotic triads, we'll use an example of *de-authentication* and proximity detectors for *computers-on-wheels (COWs)* in a hospital. One triad characterizes the creation of the security policy. The administrator perceives a reality (*referent*) where clinicians are walking away from logged-in sessions, and thus creating data exposure and corruption risk. The administrator then constructs a mental model (*thought*) where COWs automatically log out sessions when users walk away. Deciding that this is a better reality, the administrator crafts an IT configuration (*symbol*) intended to implement this policy—in this case, by installing a proximity detector on each COW, and choosing a timeout threshold threshold (triggered by the clinician moving away from the COW) after which lack of proximity effects the logout.

However, the hospital IT system has another set of actors: the clinicians who are the users. The triad involving IT configuration, user, and reality then characterizes the emergence of the workaround. The admin's new IT configuration (*symbol* generates a *reality* (*referent*) where proximity detectors cause appropriate logouts. However, the clinicians perceive this reality as not matching their desired workflow (*thought*)), where clinicians often must walk away from the COW to examine a patient, to observe readings on a device, to find a document, and to speak with another clinician. Consequently, the clinicians generate their own addition (*symbol*) to the IT configuration—inverted styrofoam cups placed over the detectors that defeat their function—to modify the generated reality (*referent*) to one closer to their liking and need. Furthermore, unless the administrator “closes the loop” by observing the disabled proximity detectors, the administrator

may never realize that the eventual result of this security improvement (automatic timeout) is an *increase* in exposure, because previously timed logouts are now indefinitely postponed.

The semiotics of language and the effective communication of meaning focus on *morphisms*—“structure-preserving mappings”—between nodes of the triad. However, with IT usability problems we are concerned instead with ineffective communication—and hence focus on what we call *mismorphisms*: mappings that *fail* to preserve important structure when we go from z in one node of the triad to its corresponding z' in another (Figure 1-b,c). Indeed, we hypothesize that mismorphisms lie at the heart of circumvention, because they characterize the scenarios that frustrate users—and often the resulting circumvention itself.

The styrofoam cup scenario above provides examples of several types: the reality generated by the new IT configuration failed to preserve the workflow features desired by the users; the administrator imagined that “dialing up” security configuration—by adding a timeout—would *increase* security; but when mapped to reality, the change *decreased* security; the users’ additions of styrofoam cups caused the emergent reality to lose the security properties the administrators imagined.

In the following sections, we explore this idea by identifying specific categories of mismorphism, and showing how they are empirically supported by items in our corpus.

4 Loss of Static Properties

Many troublesome scenarios arise when a mapping from one triad node to another fails to preserve some critical property. For clarity of presentation, we can usually treat this property as some Boolean predicate. More precisely, when z in one node of the triad maps to z' in another, we may have $\Phi(z) = \text{true}$ but $\Phi(z') = \text{false}$, for some crucial predicate Φ . (E.g., see Figure 1-d.)

4.1 Lost Workflow Properties

In many common incarnations of this type of mismorphism, the reality generated by the administrator’s IT configuration does not match the workflow the users perceive as necessary. E.g.:

- Sawchuk et al. (2014) warns against *electronic health records (EHRs)* that list test results in chronological order, causing clinicians to miss the most recent test results, which they expect to see at the front.

Here, the predicate Φ means the results are displayed in reverse chronological order; it holds in the clinician’s mental model but not in the generated reality.

- Harrison et al. (2007) lament how computerizing the prescription order entry process led to “imposition” of a “linear workflow” which ‘led to delays in delivery of orders, left nurses in doubt whether physicians had initiated orders, and sometimes produced divergent printed orders: the physician’s original and the pharmacist’s modification”—in turn leading to workarounds.

Here, the generated reality loses many of the properties present in the users’ mental models of their desired workflow.

- A vendor of power grid equipment had a marketing slide showing their default password and the default passwords of all the competitors. The slide was intended to show how secure this vendor was, since they used a more secure default password. However, a deeper issue here is that access to equipment during an emergency is critical, since availability of the grid is far more important than other classical security aspects. Any scheme to replace default passwords with a stronger scheme needs to preserve this availability.

Here, we have two predicates at play: password authentication with well-known defaults generates a reality that fails to preserve the basic security properties in the administrator’s mental model; but password authentication without well-known defaults generates a reality that fails to preserve the availability required in the domain expert’s mental model.

- In a hospital’s EHR system, there’s a box for documentation of certain information. However, the box is limited to N characters. To write more, one opens another box, and then another and another. But unless the clinician knows she is about to run out of space, there’s no way to inform the reader that text continues on the next box—and the reader sees no indication that more than one box exists.

Here, the predicate the EHR IT loses is the property that the record has a well-defined indication that “here ends the record.”

- When serving as the editor for a paper in an academic journal, one can send the paper out to reviewers, or (if the paper is poor enough or fails to meet certain criteria) reject the paper immediately. When IEEE switched to a new editing portal, there was no longer a clear way to do the latter; the workaround was to first reduce the “minimum number of reviews” required down to zero.

Here, the standard editorial workflow includes summary rejection, but the generated reality loses this.

- Use of a commercial “network flow anomaly” tool at a university led to many amusing (in hindsight) false positives, such as black-holing the Computer Science web server during the first week of classes, because students were downloading their class tools.

Here, the predicate lost was the correctness of whether or not the scenario was a genuine instance of network abuse.

In a particularly ironic twist, sometimes the technology itself, in the setting in which it’s being applied, causes the mismorphism. E.g.:

- Knowledge-based authentication at one credit bureau failed for one of the authors when he was the victim of identity theft, because the bureau assumed that the information (e.g., past addresses) in their record was accurate. However, identity thieves corrupted this information, so the genuine user was not able to correctly answer questions about it. (There were similar problems with trying to correct the “current” address, since none of the choices it gave were correct.)

Here, we lose the property that “bona fide user can authenticate himself to system” precisely because this choice of authentication technology fails when the user has been the victim of identity theft.

- Similarly, in a high-profile Twitter username theft, the genuine user failed knowledge-based authentication at several sites because the thief had corrupted the data being asked about (Hiroshima, 2014)
- A colloquium speaker on his way to Dartmouth was stranded in rural NH because he stopped to have a picnic in a cellular deadzone—and the rental car required cellular access to be started. (The workaround: have a tow truck tow the car to live zone.)

Here, we lose the property that “bona fide user can authenticate himself to system” precisely because of this choice of wireless technology.

- Too much wireless coverage can also cause problems: in an intensive care unit, RFID chips were found to interfere with about 53 of 117 devices—and for 17, the interference had serious consequences.
- Sometimes, due to things like dose sizes, a nurse must obtain more of a drug than the patient’s dose. For certain sensitive drugs, it is required that a nurse have another nurse witness the disposal of the excess. However, in some medical enterprises, it is standard practice to have the witness take place when the nurse first removes the drug from the drug dispensing machine (*pyxis*)—before the dosage has been given and the excess disposed.

This choice of workflow—witness before the disposal—can not, by definition, work to document witnessed disposal.

Failure to preserve some critical property can also develop over time. For one example, we often see *citogenesis*: when some artifact of the IT causes a spurious change to reality’s representation, which then gets interpreted by all users as genuinely representing the real world. E.g:

- Medical personnel tell of *chart ghosts*: when information mistakenly gets added to a patient’s record, it becomes real. It can retroactively change many other parts of the patient’s record—and clinicians may take the multiple occurrences of this information as confirmation that it must be true.
- In a report (Kolbe, 2014), potentially apocryphal but certainly entertaining, a prankster allegedly claims to have changed the Wikipedia page for “Chicken Korma” to indicate that “azid” was an alternate name, in order to tease a friend named Azid. The association persists in online literature.

Additionally, many amusing but dangerous distortions, building over time, stem from *note bloat*: the ease of copying and pasting previous information into a record: e.g., OK vitals are recorded at 6am for a patient who died the night before. (Indeed, the Texas Ebola patient’s record is over 1000 pages long—and also indicates that he was assigned an attending physician four minutes after he had left the hospital.)

A similar time-based mismorphism, leading to bad security practices by administrators and users alike, is underestimating the current (and growing) power of the adversary. Schneier (2014) notes that in a recent cracking challenge for a “16,000-entry encrypted password file”, the “winner got 90% of [the passwords], the loser 62%—in a few hours.” Griffith and Jakobsson (2005) earlier used public information to “reduce the entropy of a mother’s maiden name from an average of close to 13 bits down to below 6.9 bits for more than a quarter of the people targeted, and down to a zero entropy (i.e. certainty of the their mothers maiden name) for a larger number of targeted individuals.... approximately 4,105,111 Texans.” Heckle (2011) quoted a clinical user noting “you would have to be real familiar [with] how to work Meditech to be able to understand anything.”

4.2 Circumvention as Compensation

When the generated IT fails to have some property the users regard as critical, a standard circumvention response is for users to customize the IT configuration to compensate.

One standard way is to add functionality. This can include all the standard ways users share credentials (thus causing the “1-1 credential-person” property in an administrator’s mental model to fail in reality): sticky notes with passwords; shared PINs; a senior professor sharing his NSF password with a staffer; Snowden allegedly getting colleagues to log in for him (Isikoff, 2014); US ICBM launch codes allegedly set to “00000000” (Nichols, 2013). We even saw one banking scenario where employees routinely used the credential of an employee appropriately authorized but deceased. Other examples include users adding new technology:

- Security officers in industry report finding employees setting up unprotected WLAN access points in order to provide themselves features prohibited by enterprise policy.
- At a defense laboratory, the classified and unclassified networks were air-gapped, but network security officer lived in fear that a user with an urgent deadline would connect a printer/etc from one side to the other side and bridge the gap. (Here, the user’s action to make the generated reality better match his desired workflow causes the reality to *stop* preserving the admin’s desired air gap.)
- At a hospital, a clinician was unhappy with the official computer app for processing medical data, so he wrote his own to grab what he really wanted from the raw packets.
- In energy trading, daily profit/loss reports require access controls because of US Securities and Exchange Commission (SEC) rules, but these reports are so useful that traders use their own scripts to get at the database directly via well-known workarounds.

Sometimes users compensate for the loss of a critical property by *removing* functionality.

- In multiple industries, infosec officers have told us that senior staffers insist on not patching their compromised machines, sometimes by disconnecting them during remediation.
- An automotive “smart key” automatically unlocks the car when it’s within a few feet, which made it hard for surfers to leave their keys near the car when surfing (because they could not take their electronic keys into the water.) To work around this, users would remove this functionality by shielding the keys in a Faraday cage made of aluminum foil (Paul and MacNaughton, 2014).

Sometimes, the technology itself removes functionality:

- Many code features (termed “optimization unstabled”) are silently removed by optimizing compilers. These may be security-relevant (Wang et al., 2013). Here, we see two kinds of mismatches: between the reality of the deployed code and the mental model of its security-aware programmer, and also between the reality of “relevant features” and the mental model of the compiler optimizer.

Alternatively, users can establish *shadow systems*, sometimes using functionality already inadvertently present in the system—and this inadvertent presence itself can sometimes be seen as a mismorphism: the administrators would probably have not allowed this pathway had they seen how it would undermine policy. E.g.:

- Clinicians started using the telephone instead of an online service because the telephone did not require a password (Heckle, 2011).
- In trading, employees perform desired exfiltration, despite data exfiltration guards, by scanning documents, turning them into images, then embedding the images in PDFs—rendering the text opaque to the online guards. (In a different industry, we noted employees screen-scraping medical images into Powerpoint, for similar reasons.)
- A user of an online meeting service could not remember the password associated with his registered email, so he enters the service instead via an unregistered identity which required no password, essentially creating a phony duplicate of himself.
- To avoid data exfiltration guards, employees send work documents to their email at home.
- University IT wants all personnel to use enterprise Sharepoint to share data. Science professors use Dropbox instead because it easily allows external users (unlike the official service) and can be configured by the users themselves (when they’re chasing deadlines after hours).
- A confidence artist allegedly scammed Apple Store stores because their system for authenticating “override codes” on rejected debit cards merely counted the number of digits in the alleged code—so any alleged code of the right length would be accepted.
- Government employees tasked to build a porn filter could not test it without violating government policy, and so went offsite; we have also interviewed other government employees in hyper-secure facilities that prevented all external communications explain about needing to work from Starbuck’s or via tunneling to a university system. In one case, employees set up their company to avoid enterprise enforcement of regulations.

4.3 Mismorphism as Circumvention

Mismorphism can be a vector of circumvention, as well as a cause. One category we see is *intentional distortion*. A human user, striving to get the IT to generate the desired real-world functionality, intentionally alters one of the triad mappings, making it less correct. Sometimes, the user aspires to later undo this distortion; sometimes, she does not. E.g.:

- In one medical scenario, one EHR prevents the doctor treating a patient with predicted risks of clotting from leaving the software until the doctor orders a blood thinner. If the patient is already on blood thinners, the double dose may kill her. The workaround is for the doctor to order the second, lethal dose, then go back into the system and cancel the original dose.

That is, to leave the EHR, the clinician must make the “EHR reflects needed dose, not lethal dose” invariant temporarily false.

- Multiple medical personnel talk of telling the IT system that an intravenous (IV) bag is smaller than it really is, so that the “almost out” alarm goes off earlier (and indeed, one of the authors recently saw this first-hand during an emergency room visit).

To further their goal of providing better patient care, the clinicians sacrifice the invariant “IT system reflects actual dosage patient has received.”

- Some smart pumps assume the patient never weighs more than 350 pounds; for overweight patients, clinicians must similarly distort the IT representation by hooking up two pumps (each allegedly serving a patient of half the actual weight), or by telling the pump it is giving a different medication that, for a legal weight patient, works out the correct drip rate.
- In hospitals, the EHR often enforces insurance rules that a patient must have a certain diagnosis before paying for certain tests. This reality loses the property (required in the doctor’s mental model) that these tests are necessary in order to determine what the diagnosis is. To compensate for this mismorphism, doctors will commonly introduce another: intentionally adding a known incorrect diagnosis into the patient’s record.

One clinician reported seeing yet a third step in this cascade: the Cerner EHR repopulated that patient’s whole life record with that spurious diagnosis—citogenesis reputedly hard to undo.

- One of us has previously lamented (Koppel, 2014) the existence of a *make me the author* function in some EHRs. This “allows the most recent user to appear as the creator of the document, even though he or she may have only updated a number or made a quick note....the legal and medical consequences of this feature are mind-boggling.”
- A university’s travel portal requires that all trips be recorded, but only permits some authorized travelers (e.g., faculty) to have accounts.

As a workaround to this mismorphism, people with no account (e.g., grad students) have their trips recorded in the account of the faculty member who supervises them.

However, this workaround itself introduces a new mismorphism, as the travel portal also has some workflow logic that assumes that all trips recorded in a user’s account are being taken by that user. Consequently, when a professor went to set up a trip from Boston to San Diego, the portal warned him that it conflicted with his trip from Mexico City to Las Vegas—even though was not aware he was going to be in either city.

Breaking the Workaround Sometimes, we see a second round of mismorphism: the IT loses the property that the workaround works. E.g.:

- At a particular university, the email system ended up giving each user many email addresses. This technology lost the property that email from the same sender would have the same sender; however, users worked around this by learning the equivalence sets of addresses for their correspondents—“I know that ‘bjones at foo.university.edu’ denotes the same person as ‘William.Jones at univesity.edu.’”

However, when that university then added identity PKI, each user’s certificate only specified that user’s canonical university.edu email address. When Bill Jones sent signed mail from his “bjones at foo.university.edu” account, Apple’s S/MIME client would identify the signature as not valid because

the email address didn't match, even though the cryptography indicates the message was in fact signed by the holder of the private key of "William.Jones at university.edu" and had not changed since, and the receiver *knows* the two addresses specify the same person. The client (at least at that time) would not tell the user what the user needed to know.

The addition of S/MIME thus broke the workaround.

- In an EHR, a doctor could not find an appropriate place to record the medication he thought was needed (missing property). But he found a box that he thought would be seen and recorded it there (workaround). However, the box was not visible to subsequent users of this record; the order was not seen, and the patient was in crisis (failure to preserve the workaround).
- One business enforced a policy via its firewall of not allowing access to Dropbox. Employees who typically worked remotely regularly used Dropbox nonetheless (missing property), because they came into the intranet via VPN, and their home client access to Dropbox did not cross the firewall (workaround). We were told that when some of these employees showed up at the office, they complained that they couldn't reach Dropbox (failure to preserve workaround).

4.4 Provisioning

When it comes to access control specifically, a particular challenge is the difficulty (e.g., Smetters and Good, 2009) of what some industries call *provisioning*: the mapping of an administrator's mental model of "correct" access control to an IT policy configuration that generates a real-world system enforcing that model. Problems with provisioning are central to many scenarios of security engineering and circumvention trouble, but these problems themselves are consequences of mismorphism: failure of the mapping between the triad nodes to preserve certain structure. E.g.:

- A figurative "greybeard" in computer security tells of giving a room full of experienced Unix sysadmins the problem of devising a scheme in Unix filesystem access controls to match a relatively simple enterprise organization model. Each sysadmin would very quickly come up with a solution. But each solution was wrong.

Even for those understanding the provisioning technology did not come up with an IT configuration that generated a reality matching their goal. item Maxion and Reeder (2005) survey many awkward consequences of inadvertent policy: users at universities and in a P2P system unintentionally making their private files world readable; political operatives inadvertently leaving confidential senate documents exposed.

- In Reeder (2008), users were asked to create policy encodings to match various real-world semantics. Depending on the ambiguity-resolution approach used by the underlying filesystem, sometimes they got it right, and sometimes they didn't.
- In another trial, Brostoff et al. (2005) found test subjects universally fail at key elements of policy creation.

The reverse mapping—from IT policy to mental model—is also problematic.

- An investment bank had "entitlement review" in which employees reviewed their privileges and gave up ones they did not think they needed—except then they had to ask for them back (Sinclair, 2013).
- Xu and Stoller (2012) note that a barrier to automated *role mining* is "interpretability": "a role produced by a role mining algorithm will be adopted by security administrators only if they can identify a reasonable interpretation of the role."

In real-world organizations (as opposed to computer security textbooks), provisioning using standard technology can be dauntingly complex: we have seen enterprises using RBAC with more roles than employees; and a university system trying to establish a PKI for its population ended up with over 200 types of individuals. This complexity leads to workarounds: in banking, we saw *copy-and-paste provisioning*: when Bob joins the group, the manager will think “his job is similar to Alice’s” and so will copy-and-paste Alice’s entitlements to Bob (Sinclair, 2013). Of course, this workaround itself introduces more mismorphism: how the manager uses the provisioning tool no longer matches how the admin imagines the manager users it.

We see similar issues when it comes to capturing other parts of an organization’s workflow requirements:

- Commercial health IT systems require customization. At one hospital, to ensure representation, each team sent someone to the customization meetings. However, each team usually sent the most junior person, most easily spared—but also least knowledgeable about the workflow requirements, and possibly insufficiently assertive.

Here, we see two levels of mismorphism. In the administrator’s mental model of the customization process, each team’s view is effectively represented; in the generated reality, it is not. Consequently, neither does the resulting system preserve the workflow requirements of the user population.

- When prescribing medication, doctors think in terms of *total dosages (TD)*, but pharmacists often think in terms of the doses available via the manufacturers (e.g., “40mg tablets.”). When one hospital recently deployed a new computer system, it called a meeting to decide how to arrange dosing. The pharmacists came but the doctors were too busy. The computerized system thus focused first on available units, leading to orders such as “100mg via 3 40mg tablets” and nurses wondering if a razor blade was provided.

Again, the development reality did not preserve the representativeness in the admin’s mental model; the generated reality did not preserve the nurses’ requirement that orders can actually be filled.

5 Loss of Functional Properties

Both the administrators officially configuring IT systems and the users unofficially reconfiguring them are practicing a form of security engineering: trying to optimize some overall property of the system by adjusting a human-settable parameter. However, this process implicitly assumes that a functional relationship exists from the parameter to the property, and that the morphisms between nodes of the triad preserve that relationship. In many circumvention scenarios, both the causes—and sometimes the negative security consequences—stem from morphisms failing to preserve this relationship.

More specifically, in questions of security design, implementation, and use, we implicitly have some numeric function \mathcal{S} taking a tunable parameter (e.g., password length) to the level of security achieved. The intention of the human is to tune the parameter x so as to maximize $\mathcal{S}(x)$. However, if the mappings across the triad nodes fail to preserve crucial properties of this x vs $\mathcal{S}(x)$ curve, unfortunate things can happen. We discuss three properties in particular.

5.1 Loss of Monotonicity

In one node, the function \mathcal{S} can be **monotonic**: for $\Delta > 0$, $\mathcal{S}(x + \Delta) > \mathcal{S}(x)$. However, when mapped into another node, we may lose monotonicity: $\mathcal{S}'(x + \Delta) < \mathcal{S}'(x)$.

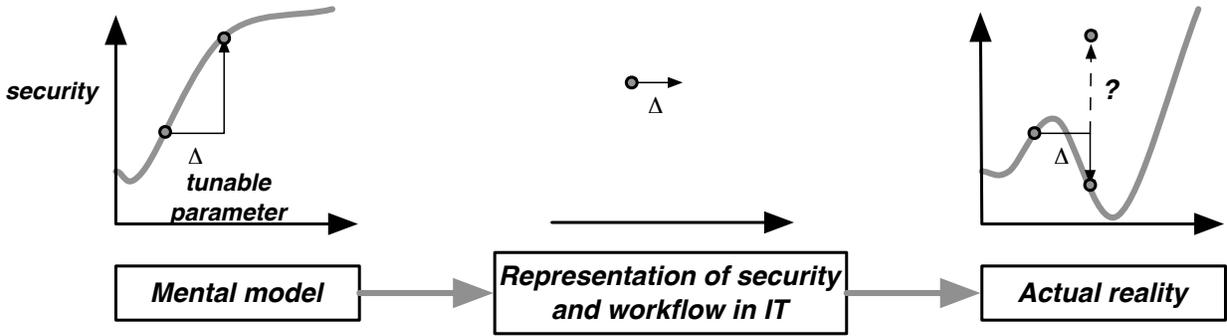


Figure 2: In what we call the uncanny descent, the mental model shows that dialing up security improves security; but when this change is mapped through the IT configuration into reality, security actually decreases.

5.1.1 Uncanny Descent

Computer graphics offers the term “uncanny valley” for when dialing up realism makes things worse before it makes things better. In security scenarios, we see many variations of such uncanniness. The time-out/styrofoam scenario is a good example of what we call *uncanny descent*: dialing up security in the IT configuration (from x to $x + \Delta$) instead leads to a decrease in security in the system itself: although the admin imagined $\mathcal{S}(x + \Delta) > \mathcal{S}(x)$, we had $\mathcal{S}'(x + \Delta) < \mathcal{S}'(x)$ instead (Figure 2).

“Best practices” for password-based authentication are notorious for exhibiting the loss of monotonicity when we go from the admin’s mental model through the IT configuration into the generated reality: dialing security up can make it worse instead. E.g., Requiring users to choose strong unique passwords for each service leads to clinicians recording passwords on sticky notes, on sticky notes provided by health IT vendors (Figure 2 in Koppel et al. (2015)), clinicians writing down passwords in notebooks (Sinclair, 2013) password notebooks for sale on eBay—and even Schneier (2005) endorsing the practice. In one defense-related industry, users reported having to change their passwords every 90 days to something new—however, going to the help desk on the 91st day and claiming to forget the new password let them restore the old one. when one of our universities established requirements that made passwords very hard to remember, many users reported relying instead on regularly resetting it via security questions that were easy to guess. A non-security graduate student observed on Facebook “everyone who is reading this already knows enough to log in as me.”

Beyond passwords, we see many other examples of dialing up a control or adding additional features counter-intuitively making things worse.

- Studies of medical workflow repeated show that adding computerized controls can impair clinician communication (Harrison et al., 2007), cause patients to disregard clinicians (Shaffer et al., 2013), cause medication errors, delay tests, delay surgery and make easier for the attending to hit one button and sign off on everyone without the irksome necessity of actually reading the individual reports.
- A study of authenticating email in power-grid emergencies found that adding S/MIME digital signatures led to users making worse decisions than they did with plaintext (Masone, 2008).
- A grade school cafeteria deployed a fingerprint reader to streamline identification and authentication of children (for lunch charges) instead of relying on often-lost ID cards. Instead, these became a bottleneck: cashiers had to constantly clean the readers of grease from French fries and such.
- An executive from a major software vendor spoke of shipping products with security controls off by default—since they were sufficiently effective at blocking behavior that users would tie up the help desk.

- A college limited the size of messages sent through the internal secure email system—leading to an employee sending a file containing over 35K personal records to a private email address, which was unfortunately mistyped, leading to inadvertent exfiltration.
- A state university established a policy that faculty email were subject to “Freedom of Information Act” requests—leading to many faculty systematically using non-university email for all their university business.
- Harrison et al. (2007) note that “when nurses under heavy work loads encounter cumbersome software requiring multiple screens for medication administration, the nurses often delay medication charting until the end of their shifts”—leading to many negative consequences.

One colleague noted the following announcement at NIST and wondered: what actual positive good was achieved by dialing up identity regulations?

New Visitor Access Requirement: Effective July 21, 2014, Under the REAL ID Act of 2005, agencies, including NIST, can only accept a state-issued driver’s license or identification card for access to federal facilities if issued by states that are REAL ID compliant or have an extension. Driver’s licenses from nine states and territories are not compliant and will not be accepted as identification: Alaska, Arizona, Montana, Oklahoma, Louisiana, Kentucky, Massachusetts, Maine, and American Samoa. In addition, NIST will accept only enhanced driver’s licenses (identified by the American Flag on the face of the card) from three states: Minnesota, New York, and Washington State

We’ve also seen incidents of faux uncanny descent: dialing up security led to an incorrect perception that actual security decreased. For example, changing an EMR to make it easier for clinicians to record when medications were given at the wrong time led to an increase in the reports of mistimed medications—which managers interpreted as a decrease in quality of service. For another example, a study of password managers (Chiasson et al., 2006) observed that when the “manager failed to generate a password considered ‘strong’ by the web site, users expressed their concern and disapproval.”

5.1.2 Uncanny Ascent

A different consequence of the loss of monotonicity is what we call *uncanny ascent*: dialing *down* the security controls can also counter-intuitively lead to an *increase* in actual security. Once again, the map from the administrator’s mental model through the IT to reality does not preserve the shape of the setting/security curve. Two examples:

- An infosec officer for a large pharmaceutical reported a nice example of this. Concerned that senior executives were illicitly sharing their work account passwords with assistants and staff, he instituted a rule that executives use the same password for both their work accounts and their personal salary and benefit information. Eliminating unique passwords (which is “bad”) led to a reduction in sharing (which is “good”). $\mathcal{S}(x - \Delta) < \mathcal{S}(x)$ but in fact $\mathcal{S}'(x - \Delta) > \mathcal{S}'(x)$.
- A common belief is that making passwords longer makes them more secure. However, a student exploring gmail’s password strength meter discovered that it considered “*qwertyqwerty* to be a weak password, *qwertyqwert* to be a fair password, and *qwertyqwer* to be a good password.” Shortening the password made Google consider it to be stronger. (We’ve also found sequences of lengthened passwords that change from strong to good to strong to good to weak—the assumed monotonic curve can in fact become rather bumpy!)
- One is not supposed to have one’s browser remember critical passwords, such as for the master employer SSO. However, a sophisticated phishing attack lured users to a rather convincing spoofed login page

for the SSO page at one of our universities; Safari tipped the user off by refusing to auto-populate the userid and password for this page because it didn't recognize it. The bad practice of having the browser remember the password protected exposure of the password to the phishing site.

- In a non-password context, Maxion and Reeder (2005) observed that “Microsoft publishes a list of ‘best practices’ for NTFS security that advises users not to use several of the provided features.”

5.1.3 Uncanny Nop

Dialing up security can also lead to no change in actual security.

- One very large IT firm required that users regularly change passwords to something distinct from the last N they had. However, as it only saved and compared against password hashes, users simultaneously complied and defeated the mechanism by changing one character.
- Multiple universities offer two flavors of WiFi—a constrained one for “public/guest” and a more powerful and privileged one for authenticated insiders—but find that in practice, massive numbers of insiders use the public/guest one, because it’s easier.
- We have seen hospitals that, on public displays, list patient names only as the first three letters and the last three letters, in order to provide privacy. However, often, a patient’s name can be sufficiently unique that this code uniquely identifies him or her to non-clinicians—thus violating privacy anyway. (Lack of uniqueness can also be risk: multiple patients may map to the same displayed identifier.)
- Enterprises intending to delete material from the Web often merely remove a published link; their servers will still happily provide the material to any browser that asks with the right URL. (Two notable examples include Texas attempting to remove death indices (Griffith and Jakobsson, 2005) and ApplyYourself trying to delay announcement of business school admission decisions (Smith, 2005).)
- Researchers repeatedly observe how educating users about good security practices does not seem to change their actual behavior (e.g., Riley, 2006; Yan et al., 2005; Dhamija and Perrig, 2000; Heckle, 2011).
- A neuro intensive care unit (ICU) added technology to track nurses with RFID tags and display their location on a central monitor. However, processing lag led to the monitor consistently displaying instead where each nurse had recently been.

In the converse, Schechter et al. (2007) found users did not notice when the *https* indicator on a Web page was removed.

5.2 Loss of Continuity

Another relevant mathematical property of functions is *continuity*. In one node, the function \mathcal{S} may be **continuous**: if $|x - y| < \delta$, then $|\mathcal{S}(x) - \mathcal{S}(y)| < \epsilon$, for small positive δ and ϵ . However, when mapped into another node, the function may lose continuity: the difference $|\mathcal{S}'(x) - \mathcal{S}'(y)|$ may be significantly large.

Circumvention scenarios often arise because the morphisms across the triad nodes fail to be continuous. Small changes in parameters lead to big changes in \mathcal{S}' in the generated reality. E.g.:

- Amusingly tangible examples here are the regular occurrences of when an innocuous photo reveals a password which users have posted on paper—the small change of a photo yields to a dramatic change in who can authenticate. Were we cavalier about copyright, we could show examples from the World Cup, the Super Bowl, and Prince William’s military service.

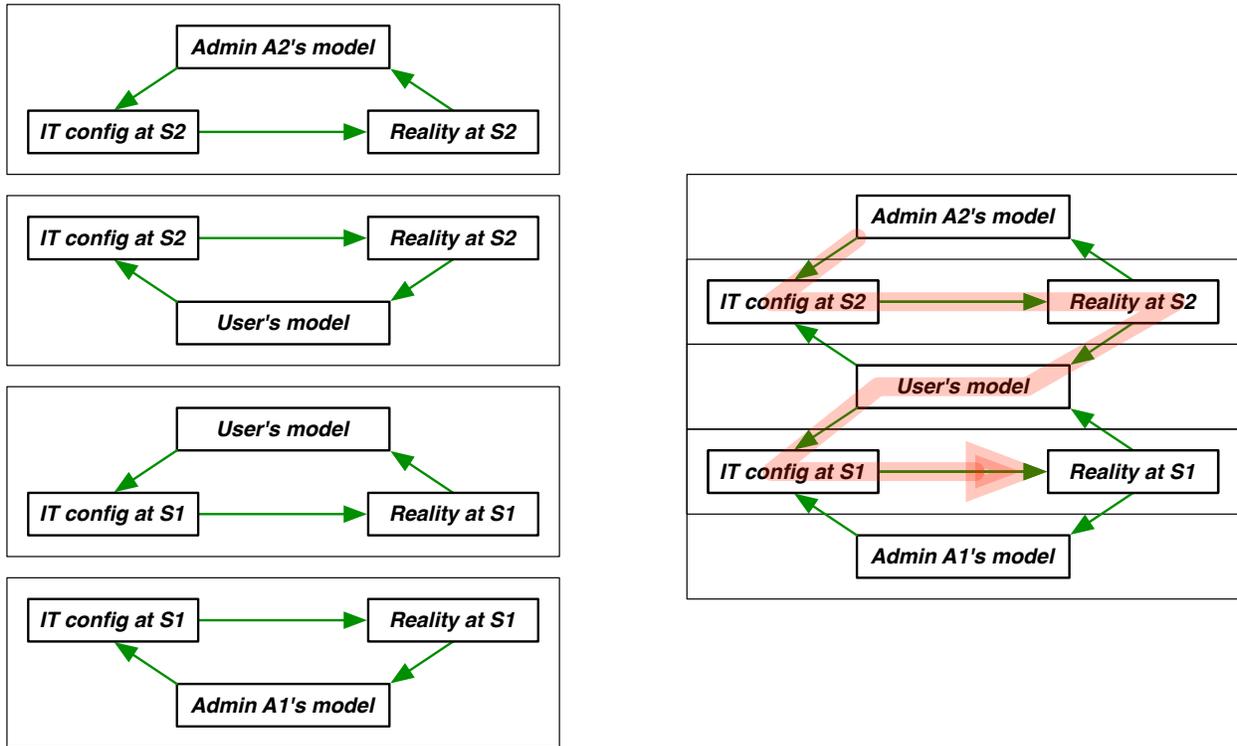


Figure 3: Although naively we might think of each admin/user as part of distinct triad (left), they can share elements. This overlap creates the potential for cascading effects (right). E.g, administrator A_1 may think his system’s security depends primarily on the configuration he chooses; however, if the administrator A_2 of a different system trains his users to accept some bad practice, then that can also affect what happens in practice at S_1 .

- A colleague reports that his personal domain name differs by one character from a large University hospital’s; in one 24-hour period, he received a hospital bulletin, an equipment invoice, an employment contract, and photos of a rectal polyp.
- GUIs can have problems when a button on one screen has the same location but very different semantics from a button on the next screen. Inadvertent clicking has contributed to tornado sirens going off by mistake—and to a patient dying due to lack of medical treatment.
- In medical device interfaces where users enter values digit-by-digit on a keypad, a small typing error (e.g., an extra zero) can lead to a significant medication error (e.g., a factor of 10 too large). Similar problems arose when the mode display quietly changed from “mg” to “mg/kg” while keeping numbers the same.

5.3 Domain and Range Trouble

Another property that can be lost to mismorphism is the nature of \mathcal{S} as a function. In a mental model, $\mathcal{S} : D \rightarrow R$ can be a well-defined function taking some $x \in D$ to $\mathcal{S}(x) \in R$. However, when we map to the generated reality, we lose these properties. Instead, perhaps \mathcal{S}' depends on other parameters besides $x \in D$; or perhaps changing to x to $x + \delta$ changes more than just items in R .

In the former case, we lose the morphological property of *locality of control*. The administrator A_1 of system S_1 implicitly assumes that de facto security of S_1 depends only on the de jure configuration A_1 puts

together—or, at worst, also on the behavior of the users. A user U of a system may believe his actions only affect his portion of the system, and not those of other users. However, given that the same user can use multiple systems, and that the same system can be used by multiple users, effects of actions can reach unexpectedly far. Such *cross-channel effects* between apparently unrelated nodes can both lead to, as well as exacerbate, the consequences of circumvention.

For example, we often see “action at a distance”—when the security of a system S_1 is reduced because of the actions of an administrator A_2 of a different system S_2 . E.g.:

- An energy trader set up an SSL server, for security, but used a self-signed certificate—thus leaving his own service vulnerable to man-in-the-middle attacks, but also training all his users to accept self-signed certificates on SSL sessions, thus increasing the exposure of all the other SSL services—banks, credit cards, medical sites—they use. The security of these other sites, in practice, decreased because of the actions of a careless administrator on an apparently unrelated site. Without changing their own x , $\mathcal{S}'(x)$ suddenly declines.
- A university requires that users enter their master SSO password into a standard http “basic authentication” pop-up, in order to access the official Sharepoint service. Here, the local administrators apparently felt that since the service could not be reached except via a protected LAN, that no exposure was created—except that by doing so, they trained their users to enter their master university SSO password (and any other password, for that matter) into a basic authentication popup anyone provides. The official Sharepoint installation might be OK, but other sites using that master password or any other password these users wield have suddenly have \mathcal{S}' drop.
- In a medical setting (Heckle, 2011), administrators added a way to quickly logout via a short keystroke sequence—thus reducing de-authentication risk on those systems, but *increasing* it on the other systems which did not support this shortcut, although users were trying it there.

However, some have observed that cross-channel effects can help: if a large Web site insists on good user password practices, then users may apply those principles elsewhere as well (Goodin, 2014).

Password practices at one site—good and bad—can create risks at other sites. In multiple studies, users report that when they are forced to change a password at one site, they will attempt to change their passwords at other sites to match (e.g., Dhamija and Perrig, 2000; Goodin, 2014). Studies also show users having the same number of unique passwords even though the number of accounts increases (Gaw and Felten, 2006) and choosing sets of passwords from the same close family (Chiasson et al., 2006; Florencio and Herley, 2007; Florêncio et al., 2014). As has been widely observed, reuse of a password across two sites means exposure at site S_1 —e.g., due to failure to protect the password file—increases the exposure risk at otherwise unconnected sites. An even more direct risk is that a user who finds a password rejected will often iterate over all the passwords in their standard set—thus providing their full set to a Trojan that rejects everything.

In the other direction ($\mathcal{S}'(x)$ having more components than allowed for in $\mathcal{S}(x)$), a basic pattern occurs when users choose actions which they perceive as maximizing their own utility, but which hurts all of a system’s users—such as the university staffers whose seniority lets them persist with one-character passwords, or the investment traders who disconnect their systems during remediation so they remain unpatched. (Essentially, this is the tragedy of the commons.)

6 Related Work

The classic work of Ogden and Richards (1927) generated some subsequent scholarship relevant to computer systems including the use of formal semiotic models to examine user interfaces and human access to the underlying computational functionality (e.g., Ferreira et al., 2005; Goguen, 1999; de Souza et al., 2001).

More recently, several researchers have investigated the effects of user mental models on their security decision-making (e.g. Wash, 2010; Camp, 2009; Olembo et al., 2013). An understanding of mental models may help predict human behavior that would otherwise seem irrational but is rational in the context of a faulty model (Johnson-Laird, 1986), and some of our work has been aimed at simulating mental models for prediction (Blythe and Camp, 2012).

Our earlier work explored workaround in bar-coded medicine administration systems (Koppel et al., 2008) and in access control issues in real-world enterprises (Sinclair and Smith, 2010; Sinclair, 2013). The short paper Blythe et al. (2013) surveyed our initial look at circumvention in general; the subsequent Koppel et al. (2015) looks at circumvention in medical IT. Also as noted earlier, our 2014 *JAMIA* paper (Smith and Koppel, 2014) exploring usability problems in health IT inadvertently rediscovered essentially the same structure as Ogden and Richards; we expand that model here using a much larger corpus than in our initial circumvention papers. We have also been exploring the use of agent-based modeling to help clarify the security engineering decisions involved in IT prone to circumvention.

7 Conclusion

We have presented our model looking at computer/workflow usage as an Ogden-Richards semiotic triad, but considering instead how the mappings fail to preserve structure: static properties, correspondence of “security setting” to “security achieved,” continuity, control. To support this model, we cite many examples of distortions and unwanted effects arising from mismorphisms among users’ needs, computer-embedded rules, and the (mis)understandings of computer system administrators. Such problems exist in almost every interaction with a computer system. Readers need only recall their most recent efforts at completing a web-based input form that asks the logical equivalent of: “do you walk to work or do you carry your lunch?” When these conundrums are elevated to prohibiting legitimate users’ access to needed computer systems, or preventing key functions of the organizational mission, they become irksome and they generate endemic circumvention. We offer a typology of these problems, and also explore users’ creativity in working around the pandemic challenges they create.

We emphasize that these examples do not reflect the work of those intent on misusing the system for malicious reasons. These are just users trying to do their jobs. However, many of the workarounds create vulnerabilities that are usually unknown to their creators and unknown to the system administrators who designed the systems that necessitate the circumventions.

Building this topology also highlights the necessity for observation of use in reality, rather than as reflected in the system’s blueprint or initial design.

In future work, we plan to distill this model into design principles for better security engineering. One may start by looking at mismatches as while moving around the triad and then considering where “shape” fails to be preserved, perhaps via feedback loops, regular discussions, and explicit monitoring. Alternatively, growing this corpus may allow us to create a database that security personnel can consult for design patterns. Discovering circumventions and analyzing their causes can improve system design so that users can get their jobs done without working around the rules.

Acknowledgment

This material is based in part upon work supported by the Army Research Office under Award No. W911NF-13-1-0086. The views and conclusions are of the authors alone.

References

- Bernard, H. R. and Ryan, G. W. (2010). *Analyzing Qualitative Data: Systematic Approaches*. Sage Publications.
- Blythe, J. and Camp, L. J. (2012). Implementing mental models. In *Security and Privacy Workshops (SPW), 2012 IEEE Symposium on*, pages 86–90. IEEE.
- Blythe, J., Koppel, R., and Smith, S. W. (2013). Circumvention of security: Good users do bad things. *Security & Privacy, IEEE*, 11(5):80–83.
- Brostoff, S., Sasse, M., Chadwick, D., et al. (2005). R-What? Development of a Role-Based Access Control (RBAC) Policy-Writing Tool for e-Scientists. *Software Practice and Experience*, 35(9):835–856.
- Camp, L. (2009). Mental models of privacy and security. *IEEE Technology And Society Magazine*, 28(3).
- Charmaz, K. (2003). Grounded theory. In *The SAGE Encyclopedia of Social Science Research Methods*, pages 440–444.
- Chiasson, S., van Oorschot, P., and Biddle, R. (2006). A usability study and critique of two password managers. In *Proceedings of the 15th USENIX Security Symposium*.
- de Souza, C. S., Barbosa, S. D. J., and Prates, R. O. (2001). A semiotic engineering approach to user interface design. *Knowledge-Based Systems*, 14(8):461–465.
- Dhamija, R. and Perrig, A. (2000). Deja Vu: A User Study Using Images for Authentication. In *Proceedings of the 9th USENIX Security Symposium*.
- Ferreira, J., Barr, P., and Noble, J. (2005). The semiotics of user interface redesign. In *Proceedings of the Sixth Australasian Conference on User Interface - Volume 40, AUIC '05*, pages 47–53.
- Florencio, D. and Herley, C. (2007). A large-scale study of web password habits. In *Proceedings of the 16th International Conference on World Wide Web*, pages 657–666. ACM.
- Florêncio, D., Herley, C., and Van Oorschot, P. C. (2014). Password portfolios and the finite-effort user: Sustainably managing large numbers of accounts. In *Proc. USENIX Security*.
- Gaw, S. and Felten, E. W. (2006). Password management strategies for online accounts. In *Proceedings of the Second Symposium on Usable Privacy and Security*, pages 44–55. ACM.
- Goguen, J. (1999). An introduction to algebraic semiotics, with application to user interface design. In Nehaniv, C. L., editor, *Computation for Metaphors, Analogy, and Agents*, pages 242–291. Springer-Verlag.
- Goodin, D. (2014). Apple.com does more to protect your password, study of top 100 sites finds. *ArsTechnica*.
- Griffith, V. and Jakobsson, M. (2005). Messin’ with Texas: Deriving Mother’s Maiden Names Using Public Records. In *Applied Cryptography and Network Security (LNCS 3531)*, pages 91–103.
- Harrison, M., Koppel, R., and Bar-Lev, S. (2007). Unintended consequences of information technologies in health care—an interactive sociotechnical analysis. *Journal of the American Medical Informatics Association*, 14:542–549.
- Heckle, R. (2011). Security Dilemma: Healthcare Clinicians at Work. *IEEE Security and Privacy*, 9(6):14–19.
- Hiroshima, N. (2014). How i lost my \$50,000 twitter username. *Gizmodo*.
- Isikoff, M. (2014). Exclusive: Snowden Swiped Password From NSA Coworker. *www.nbcnews.com*.
- Johnson-Laird, P. (1986). *Mental models: towards a cognitive science of language, inference, and consciousness*. Harvard University Press.
- Kolbe, A. (2014). How pranks, hoaxes and manipulation undermine the reliability of Wikipedia. *Wikipediocracy*.
- Koppel, R. (2014). Illusions and delusions of cut, pasted, and cloned notes: ephemeral reality and pixel prevarications. *CHEST*, 145:444–445.

- Koppel, R., Smith, S., Blythe, J., and Kothari, V. (2015). Workarounds to Computer Access in Healthcare Organizations: You Want My Password or a Dead Patient? In *Information Technology and Communications in Health (ITCH)*. (to appear).
- Koppel, R., Wetterneck, T., Telles, J. L., and Karsh, B.-T. (2008). Workarounds to barcode medication administration systems: their occurrences, causes, and threats to patient safety. *Journal of the American Medical Informatics Association*, 15(4):408–423.
- Masone, C. (2008). *Attribute-Based, Usefully Secure Email*. PhD thesis, Department of Computer Science, Dartmouth College.
- Maxion, R. and Reeder, R. (2005). Improving User-Interface Dependability Through Mitigation of Human Error. *International Journal of Human-Computer Studies*, 63(1):25–50.
- Nichols, T. (2013). UPDATE: Were U.S. nuclear codes set to zero? Bruce Blair responds. *The War Room: Tom Nichols on Politics and Foreign Policy*.
- Ogden, C. and Richards, I. (1927). *The Meaning of Meaning*. Harcourt, Brace and Company.
- Olembo, M. M., Bartsch, S., and Volkamer, M. (2013). Mental models of verifiability in voting. In Heather, J., Schneider, S., and Teague, V., editors, *E-Voting and Identify*, volume 7985 of *Lecture Notes in CoMputer Science*, pages 142–155. Springer Berlin Heidelberg.
- Paul, C. and MacNaughton, W. (2014). Smart key, pretty dumb. *techpageone.dell.com*.
- Pettigrew, S. F. (2000). Ethnography and Grounded Theory: a Happy Marriage? In *NA - Advances in Consumer Research*, volume 27, pages 256 – 260. Association for Consumer Research.
- Reeder, R. (2008). *Expandable Grids: A user interface visualization technique and a policy semantics to support fast, accurate security and privacy policy authoring*. PhD thesis, School of Computer Science, Carnegie Mellon University.
- Riley, S. (2006). Password security: What users know and what they actually do. *Usability News*, 8(1):2833–2836.
- Sawchuk, M. et al. (2014). *The Essential Role of Laboratory Professionals: Ensuring the Safety and Effectiveness of Laboratory Data in Electronic Health Record Systems*. Center for Surveillance, Epidemiology and Laboratory Services, Centers for Disease Control and Prevention.
- Schechter, S., Dhamija, R., Ozment, A., and Fischer, I. (2007). The Emperor’s New Security Indicators. In *IEEE Symposium on Security and Privacy*, pages 51–65.
- Schneier, B. (2005). Write down your password. *Schneier on Security*.
- Schneier, B. (2014). Choosing secure passwords. *Schneier on Security*.
- Shaffer, V. et al. (2013). Why do patients derogate physicians who use a computer-based diagnostic support system? *Medical Decision Making*, 33.
- Sinclair, S. (2013). *Access Control In and For the Real World*. PhD thesis, Department of Computer Science, Dartmouth College.
- Sinclair, S. and Smith, S. (2010). What’s Wrong with Access Control in the Real World. *IEEE Security and Privacy*, 8(4):74–77.
- Smetters, D. and Good, N. (2009). How Users Use Access Control. In *Proceedings of the 5th Symposium on Usable Privacy and Security*.
- Smith, S. (2005). Pretending that Systems are Secure. *IEEE Security and Privacy*, 3(6):73–76.
- Smith, S. W. and Koppel, R. (2014). Healthcare information technology’s relativity problems: a typology of how patients? physical reality, clinicians? mental models, and healthcare information technology differ. *Journal of the American Medical Informatics Association*, 21:117–131.
- Wang, X., Zeldovich, N., Kaashoek, M., and Solar-Lezama, A. (2013). Towards Optimization-Safe Systems: Analyzing the Impact of Undefined Behavior. In *Proceedings of the 24th ACM Symposium on Operating Systems Principles*.

Wash, R. (2010). Folk models of home computer security. In *Proc Symposium on Usable Privacy and Security*.

Xu, Z. and Stoller, S. (2012). Algorithms for mining meaningful roles. In *Proceedings of the 17th ACM Symposium on Access Control Models and Technologies*, pages 57–66.

Yan, J. et al. (2005). The memorability and security of passwords. In Cranor, L. and Garfinkel, S., editors, *Security and Usability*. O'Reilly.