

Mismorphism: a Semiotic Model of Computer Security Circumvention

S. Smith¹, R. Koppel², J. Blythe³ and V. Kothari¹

¹Dartmouth College, United States

²University of Pennsylvania, , United States

³Information Sciences Institute, University of South California

e-mail: sws@cs.dartmouth.edu

Abstract

In real world domains, from healthcare to power to finance, computer systems are deployed with the intention of streamlining and improving 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. Well-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. This paper supports this claim with empirical data from the corpus.

Keywords

Circumvention, authentication, authorization, usability.

1. Introduction

Users systematically work around security controls. The security community can pretend this does not happen, but it does. This paper reports on research addressing 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, this approach instead observes and records what really goes on compared to the various expectations. Then, after data items are reviewed, structure and models are developed, and additional data is brought in to support, reject, and refine these models. Over the last several years, via interviews, observations, surveys, and literature searches, the authors have explored the often-tenuous relationship among computer rules, users' needs, and designers' goals of computer systems. A corpus of hundreds of circumvention and unusability scenarios has been collected and analyzed. This corpus cataloged close to 300 examples of these "misunderstandings" and the circumventions users undertook to accomplish their needed tasks. The examples were derived from 285 different sources and categorized into 60 fine-grained codes. Because several examples reflect multiple codes, there were 646 applications of the codes linked to the examples.

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. Our research suggests that these triads provide a framework to illuminate, organize, and analyze circumvention problems.

This paper presents these ideas and supports them with examples. Our longer technical report (Smith et al., 2015) provides a far more exhaustive presentation of examples. When this paper does not cite a source, the example came from interviews with parties who wish to remain anonymous. As this research focuses on developing a typology rather than supporting a hypothesis, many of the usual factors in confirmation bias do not apply.

2. A Semiotic Model for IT Usability Trouble

Our previous paper (Smith and Koppel, 2014), organizing an earlier corpus of usability problems in health IT into a coherent typology, 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.

However, this tripartite framework goes back 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; and 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 words, he thinks of “Canada”—and (besides not being Mexico) his imagined Canada may differ substantially from the real one.

As our research moves from health IT usability to consider a new corpus of scenarios in security circumvention and other authentication misadventures, this framework also applies. Each scenario has at least one IT system. 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, "policy" refers to the actual machine-actionable expression of administrator 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:

- *Referent* \rightarrow *thought*: the administrator constructs a mental model of what she imagines are the actual enterprise workflow requirements.
- *Thought* \rightarrow *symbol*: the administrator 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, there now exists 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).

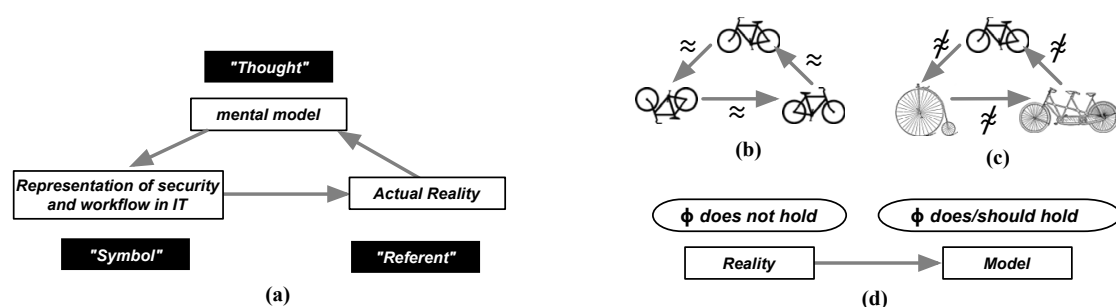


Figure 1: (a) The basic Ogden-Richards triad, moved into 21st-century IT; the arrows indicate the main direction of mappings. (b) Standard semiotics considers structure-preserving mappings between the nodes of the triad; (c) circumvention semiotics considers mappings that fail to preserve structure.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

3. Extending this Model to Security Circumvention

To illustrate the role of semiotic triads, consider 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 (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 administrator'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 one is concerned instead with ineffective communication—and hence focus on what our research calls *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, one may 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.

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, this paper will treat this property as some Boolean predicate. More precisely, when z in one node of the triad maps to z' in another, it may be that $\Phi(z) = \text{true}$ but $\Phi(z') = \text{false}$, for some crucial predicate Φ . (E.g., see Figure 1-d.)

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., 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 must preserve this availability. Here, two predicates are 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 particularly ironic twist, sometimes the technology itself, in the setting in which it is 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, the mapping loses the property that "the 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.

Failure to preserve some critical property can also develop over time. For one example, one often sees *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.

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; in one banking scenario, employees routinely used the credential of an employee appropriately authorized but deceased. Sometimes users compensate for the loss of a critical property by *removing* functionality: in multiple industries, security officers have told us that senior staffers insist on not patching their compromised machines, sometimes by disconnecting them during remediation.

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.: 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 employees screen-scraped medical images into Powerpoint, for similar reasons.)

Mismorphism as Circumvention Mismorphism can be a vector of circumvention, as well as a cause. One category seen 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.

Breaking the Workaround Sometimes, there exists a second round of mismorphisms: the IT loses the property that the workaround works. E.g. 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).

Provisioning When it comes to access control specifically, a particular challenge is the difficulty 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 system administrators the problem of devising a scheme in

Unix filesystem access controls to match a relatively simple enterprise organization model. Each system administrator 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.

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). In real-world organizations (as opposed to computer security textbooks), provisioning using standard technology can be dauntingly complex.




Figure 2: In *uncanny descent*, the mental model shows dialing up security improves security; but when mapped into reality, security actually decreases.

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, there implicitly exists some function 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 $S(x)$. However, if the mappings across the triad nodes fail to preserve crucial properties of this x versus $S(x)$ curve, unfortunate things can happen. This paper discusses three properties in particular.

Loss of Monotonicity In one node, the function S can be *monotonic*: for $\Delta > 0$, $S(x+\Delta) > S(x)$. However, when mapped into another node, the function loses monotonicity: $S'(x+\Delta) < S'(x)$.

Computer graphics offers the term “uncanny valley” for when dialing up realism makes things worse before it makes things better. Security scenarios show many variations

of such uncanniness. The timeout/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 administrator imagined $S(x+\Delta) > S(x)$, the reality has $S(x+\Delta) < S(x)$ instead (Figure 2).

“Best practices” for password-based authentication are notorious for exhibiting the loss of monotonicity when going from the administrator’s mental model through the IT configuration into the generated reality: dialing security up can make it worse instead. E.g., when one of our universities established requirements that made passwords hard to remember, many users reported relying instead on regularly resetting it via security questions that were easy to guess. (Our full report discusses many other examples.)

Fieldwork also reveals 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.

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:

- A security 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”). $S(x-\Delta) < S(x)$ but in fact $S'(x-\Delta) > S'(x)$.
- 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. (Our work has 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!)

Loss of Continuity In one node, the function S may be *continuous*: for any $\epsilon > 0$, there is a $\delta > 0$ so if $|x - y| < \delta$, then $|S(x) - S(y)| < \epsilon$. When mapped into another node, the function may lose continuity: the difference $|S'(x) - S'(y)|$ may be significantly large.

Circumvention scenarios often arise because the morphisms across the triad nodes fail to be continuous. 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.

Domain and Range Trouble Another property that can be lost to mismorphism is the nature of S as a function. In a mental model, $S : D \rightarrow R$ can be a well-defined function taking some x in D to $S(x)$ in R . However, mapping to the generated reality, loses these properties. Instead, perhaps S' depends on other parameters besides x' in D' ; or perhaps changing to x' to $x' + \delta$ changes more than just items in R' , so that the mapped range loses important information.

In the former case, the mapping loses 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, one often sees “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 , $S'(x)$ suddenly declines. Password practices at one site—good and bad—can create risks at other sites, as our full report discusses.

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). Our longer technical report (Smith et al., 2015) surveys our earlier work exploring aspects of this space.

7. Conclusion

This paper has 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, this paper cited many examples of distortions and unwanted effects arising from mismorphisms among users’ needs, computer-embedded rules, and the (mis)understandings of computer system administrators; our full report catalogs many, many more examples. 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.

8. Acknowledgements

This material is based in part upon work supported by the Army Research Office under Award No. W911NF-13-1-0086.

9. References

- Bernard, H. R. and Ryan, G. W. (2010). *Analyzing Qualitative Data: Systematic Approaches*. Sage Publications.
- 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.
- 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.
- 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.
- 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.
- Johnson-Laird, P. (1986). *Mental models: towards a cognitive science of language, inference, and consciousness*. Harvard University Press.
- 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.

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.

Sinclair, S. (2013). *Access Control In and For the Real World*. PhD thesis, Department of Computer Science, Dartmouth College.

Smith, S. 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.

Smith, S., Koppel, R., Blythe, J., and Kothari, V. (2015). *Mismorphism: a Semiotic Model of Computer Security Circumvention (Extended Version)*. Computer Science Technical Report TR 2015-768, Dartmouth College. www.cs.dartmouth.edu/reports/TR2015-768.pdf

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