

# Laws of Authentication

## [Position Paper]

Sean Peisert  
 UC Davis and Berkeley Lab  
 California, USA  
[peisert@cs.ucdavis.edu](mailto:peisert@cs.ucdavis.edu)

Ed Talbot  
 UC Davis  
 California, USA  
[edward.talbot@gmail.com](mailto:edward.talbot@gmail.com)

Tom Kroeger  
 Sandia National Laboratories  
 California, USA  
[tmkroeg@sandia.gov](mailto:tmkroeg@sandia.gov)

### ABSTRACT

In the real world we do authentication hundreds of times a day with little effort and strong confidence. We believe that the digital world can and should catch up. The focus of this paper is about authentication for critical systems. Specifically, it is about the fundamentals for measurably evaluating whether nor not someone is who they say they are. We present a “gold standard” for authentication that builds from what we naturally and effortlessly do everyday in a face-to-face meeting. We present a wide range of techniques that can and should be used in combination to provide a measurable, comprehensive, and continuous authentication system for critical systems. We also consider how such authentication systems can enable resilience to users under duress. This work differs from much of the other work in authentication first by focusing on measurable authentication techniques and also by using a multifaceted approach that integrates multiple authentication factors without adding burdensome overhead to the users.

### Keywords

Authentication, biometrics, continuous authentication, laws of authentication, measurable authentication

“Who are you?” said the Caterpillar.

*Alice replied, rather shyly, “I—I hardly know, sir, just at present—at least I know who I was when I got up this morning, but I think I must have changed several times since then.”*

—Lewis Carroll, *Alice’s Adventures in Wonderland* (1865)

### 1. INTRODUCTION

Systems can be measurably secured for attacks against availability, confidentiality, and integrity using clean-

slate, ground-up techniques involving combinations of formal verification and both technological and human Byzantine fault tolerance [30]. However, all such systems—even high-assurance, critical systems that use diverse, independent redundancy from the atoms composing a system’s transistors to the humans operating the system—require measurably strong validation and verification of human identity. We seek solutions for authenticating in critical system environments.

*Authentication*, sometimes called *origin integrity*, is a means of measuring the degree of trust that one can have that the source of data is who it purports to be [6, §1.1.2]. Humans have authenticated to each other throughout history. Some of the time, that authentication has been two people physically near each other. Sometimes two people cannot be near each other, however, or may not know each other’s appearance, so alternate means have been used, such as using the impressions of signet rings in wax, secret handshakes, or passwords. The reality of these latter techniques is that they often failed. Today, we still authenticate to each other by recognizing one another when we are in close physical proximity.

But as it has always been, there are times when authentication over a distance is required. And indeed, the techniques that we use today often fail as well. We assert that this is because most of the techniques currently used for authentication over a distance do not *measurably* provide a degree of trust. Knowing a password may indicate nothing more than that the password has not been guessed, and possessing an RSA token may indicate nothing more than it has been stolen. And it is impossible to measure the risk of either.

Thus, the focus of this paper is twofold: first, it is about methods for *measurably* evaluating whether or not someone is who they say they are, and that they are not, for example, performing a masquerade attack [24] by presenting stolen, forged, duplicated, guessed, or mimicked credentials. This paper is also about making sure that the authentication is intentional and not coerced. We briefly touch on the problem of securing systems against people who are already trusted and who then decide to do something malicious (e.g., “insiders” [8]) but primarily leave system design, formal methods, and

fault tolerance as defenses to address such threats [30]. We conclude the paper by presenting four “laws of authentication” that we assert that systems must adhere to in order to measurably capture the elements that make the current gold standard of in-person, human-to-human function well.

## 2. BACKGROUND

In the physical world, we perform authentication trivially. For example, humans may see someone whose face they recognize in a context they know, such as their workplace. If they see someone who they do not know in a sensitive area, then suspicion is raised. If they see someone they know in a place they do not expect to see them (e.g., a bank in the Cayman Islands), their suspicion may also be raised. It is this trivial, intuitive function that humans perform naturally. But it is difficult to be in a situation in which humans can always access computational resources by authenticating to another human. So we rely on computers to perform this function. We note that the Turing Test [37] represents a lighter class of the authentication problem and has shown great difficulty in simply identifying that someone is a human. Identifying that someone is a specific person is a much harder problem that people are far better suited for. Authentication to a computer today typically relies on one or more of the following (called *multifactor authentication*): something you know, such as passwords; something you have, such as RSA SecurID tokens; and/or something you are, such as some form of biometric [6, §12]. Individual authentication techniques typically include traditional passwords, graphical and video passwords, cognitive (e.g., word association) passwords, “tokens” (paper, hardware, etc..), and biometrics. Individually, these authentication techniques have a wide range of user effort, scalability, learning curve, accessibility, and resilience to theft, observation, and guessing [9]. These criteria are likely all be relevant to traditional computing, but traditional computing has a vastly different set of criteria than our question of measurably validating identity in critical environments.

The state of digital authentication is not very impressive. Even for non-critical environments, one need only look at the number of times that unauthorized users access resources they should not have access to, or authorized users are accidentally granted access to resources that they should not even with these techniques, to know that these techniques are grossly inadequate. Moreover, as we discussed earlier, they do not actually address the problem: measurably helping to verify the identity and intent of a individual. For example, long, complicated passwords are useful defenses against the threat of someone stealing a file containing password hashes. However, they are no longer relevant when the system being authenticated to locks out or throttles passwords after sufficient incorrect guesses [5]. For example, Person A knowing a password does not give Person B any measurable means of determining whether Person A is who they claim to be. Even combinations of different classes of techniques—e.g., combining a pass-

word with an RSA SecurID token and word association still do not address this problem [9].

And yet, since humans perform authentication effortlessly, intuitively, and naturally we tend to take it for granted. We instinctively assign strong authentication to even the flimsiest evidence. The continued domination of passwords over other methods of end user authentication is an example of such instinctive assignment. Even though 40 years of research have demonstrated that passwords are plagued by security problems and openly hated by users, few alternatives have emerged that perform better than passwords with regards to security, usability, and deployability [9].

## 3. GOALS AND USE CASES

We assert that physical, in-person interaction between two people who recognize each other is the ‘gold standard’ for authentication. It is *not* perfect. Human memories can fail and people’s appearance can become less recognizable (e.g., due to age) or masqueraded (e.g., via surgery). However, human-to-human is based on more than a snapshot of appearance. There are gestures, gaits, and other patterns that provide substantially more input. Studies have shown that words only convey about 7% of the message content in a face-to-face exchange. The remaining 93% is conveyed through tone of voice and body language.

The strength of authentication via passwords is very limited [25]. Password authentication to computers was developed in an environment where every bit, flop, and memory cell was precious. Technology today enables far richer authentication protocols. Therefore, the standard for comparison of authentication schemes should be the canonical face-to-face encounter between humans and not the existing password paradigm. Putting this situation in context, we can develop a simple model for human interaction.

We propose a simple model for human interaction that enables consideration of authentication in technology-enabled situations. This model includes face-to-face communication, communication through a pipe (or conduit), and communication with delay (or storage). These modes of human interaction are developed in the sections below:

*Face-to-Face* communication is the degenerate case of human (two or more) interaction. Such communication is limited because it requires both geographic and temporal synchronization. Technology enables asynchronous communication but that convenience comes at a cost, i.e., eroded authentication confidence. Passwords are the de facto means of mitigating such erosion. The judicious application of technology in the other modes of communication below can enable dramatic improvements in authentication confidence through increased natural cueing. The objective in all other modes of human interaction should be to provide authentication certainty strive to be as strong as face-to-face interaction.

A *pipe (or conduit)* enables human interaction without requiring geographic synchronization. Videoconferencing and telephone conversations are examples of such pipe-enabled interaction. Both video and telephone enable authentication via tone of voice. Videoconferencing provides richer authentication because it enables the users to view geographic context and, perhaps even more importantly, body language.

*Delay (or storage)* enables human interaction in the absence of both geographic and temporal synchronization. E-mail is an example of such delay-enabled interaction (as are social networking sites such as Facebook). The fact that e-mail uses a server to store messages, requires that users authenticate themselves to a machine to use the system. The machine becomes an intermediary between humans and the interaction is now mediated by authentication that the machine is capable of handling (i.e., passwords). Richer authentication of e-mail can be enabled by eliminating the server in favor of a peer-to-peer e-mail paradigm.<sup>1</sup>

A “While You Were Out” sticky note is an example of human interaction that requires geographic synchronization without requiring temporal synchronization. Because this interaction does not include a technological component, we include consideration of this case only for the sake of completeness.

Applying our model for human interaction suggests *richer* authentication techniques—techniques that communicate substantially more data about a person *relevant to* the authentication process—than those that are commonly used today. Rich authentication demands that users reveal significant information about themselves. This is a natural outcome in face-to-face encounters, but it is suspect in technology-enabled interaction because the technology does not necessarily demand such revelation. While our focus is on critical systems and therefore largely ignores privacy, one could imagine a scenario in which applying these techniques in non-critical environment. In such cases, privacy considerations may have more weight. For simple transactions such as browsing the news on a website, users may choose to reveal very little about themselves and the news service may require only modest user information. As the transaction becomes more important (e.g., banking or national security) the user may be required to reveal more information about themselves. Trust negotiation approaches have been developed to facilitate such interaction [22].

Certainly the highest level of authentication also demands that users be the ultimate arbiters of authentication. Machines can provide information to facilitate such arbitration but the authentication decision must rest with humans and humans alone. It is using this notion that we discuss our goals and assumptions:

**Goals.** Authentication should be measurably precise. It should never be accidental [34]. It must not be sharable, or vulnerable to loss, theft, forgery, duplication, guessing, or mimicry [35]. Ideally such a system would allow for conditions where users are under duress and enable long term auditing. This will reduce vulnerability to coercion [10, 31]).

**Assumptions.** The system, including communications between remote sensors, must be measurably secure and trustable by both parties involved in the authentication. It should also tolerate the basic tenets of security including *insider threats*, be they authorized users “gone bad” or authorized users making mistakes. This can be done, for example, by authenticating several people and requiring consensus among a majority of those people for an action to take place. We assume that electronic communications are digitally signed using means that are not easily forgeable. The implementation of these assumptions are beyond the scope of this paper, however.

## 4. PREMISE OF SOLUTIONS

We now discuss the premise of some possible solutions that fit our goals and why a number of alternative solutions do not.

### 4.1 Measurable Authentication

Collectively, rich authentication technologies, including biometrics, facial, environmental, voice recognition, GPS, etc..., readily available, and can provide measurable identity that approaches that of the physical world. Many smart-phones incorporate cameras, GPS receivers, accelerometers, and rate gyros that can identify location [17], measure gestures [27], daily movements [32], and writing style [26]. Additional, trivially available, sensors such as galvanic skin response and eye tracking [21] enable continuous real-time user authentication. Early steps in rich authentication have been promising (“Cell phones show human movement predictable 93% of the time.”).

Of all of these techniques, however, the only individual class of techniques resistant to all of these threats, in principle, are biometrics. And, not only are biometrics capable of being resistant to these threats, but they are *measurably* resistant. That is to say, while there is no way to accurately predict how likely is, even in the most crucial cases, that someone will forget their password or lose their token [2, 4]. In contrast, biometrics offer distinct advantages over traditional authentication schemes because we can measurably predict the rate of false positives and false negatives based on the type of biometric used [1, 11, 18, 29].

But biometrics applied incorrectly can still be seriously flawed. For example, in a situation without a measurably secure biometric reader that the person authenticating trusts, the person authenticating risks that their biometric may be captured and replayed in the future.

---

<sup>1</sup>Everyone running their own mail server does not scale very well in terms of effort.

And, in a situation without a measurably secure biometric reader that the person or system being authenticated to trusts, the person being authenticated to risks being the target of a replay attack. Thus, both sides must have a probably trusted means of reading biometrics, such as the *clean slate* solution referred to earlier [30]. And moreover, biometrics still result in a degree of confidence, not an absolute certainty.

Thus, we return to our previous assertion that computers lack the intuition that humans can benefit from. This does not mean that computers should not be part of the authentication equation. They are very effective in data correlation and tracking and should be used where they are strong. We believe that their role should focus on providing information, not deciding what to do with that information [16]. Thus, a computer can provide this degree of confidence to a human, but ultimately, a well-trained human is best able to make the decision. For this reason, audio and video should be also communicated as a means of providing as much as possible of an in-person, human-to-human authentication as possible. For example, authentication of an email messages can be practically validated by encoding that message with a continuous video of a person typing the email (and indicating keystroke cadence) and then entering their thumbprint and retina scan, thus communicating several types of measurable biometrics in the process combined with human visual and audio cues. While one may comment that these video streams could be faked we reiterate that our assumption is that trust in the hardware is accomplished through previously-presented means [30].

Biometrics and video—even several biometrics fused together to make masquerade harder—are not sufficient to verify intent, however. For this reason, a “secret” of some kind must also be used so that intent to access can be distinguished from accessing under duress [10, 31]. Such a secret could be a password, but a password to disambiguate intent from duress need not be one that can withstand months of brute force attempts to guess the password. That is not the objective here. The *biometrics* are the system used to provide measurable confidence of identity. The secret or password must simply enable communication of the users intent in a way that would be unlikely to be easily guessed [33]. In theory, such a secret could even be the knowledge of which finger to use to authenticate, with one finger indicating legitimate intent and another indicating duress. The space of finger combinations is probably too small to reasonably prevent guessing, however. But even a simple three-digit code<sup>2</sup> is unlikely to be guessed in the time that intent is communicated. Such a code need not have combinations of digits, punctuations, and upper and lower case letters, and need not be changed ever six months. Nevertheless note that such fingers and secrets can be vulnerable to unauthorized disclosure which is why they must be used in conjunction

with other techniques and why appropriate safeguards are still important.

“Usability” is not central to our theme but it is noteworthy. Currently, we have defaulted to a rather limited condition where we mistake burdensome security for good security: the assumption is that the more burdensome security measures are on authorized users, the more secure we are against unauthorized users. We’re finding that this is not the case [15], and is often counterproductive [13]. However, our proposed solution not only provides means to measure how likely someone is who they claim to be, but does so with less burden than existing means. Someone wishing to authenticate with our scheme could literally walk into a room naked, carrying nothing, remembering virtually nothing, and can still authenticate.

Finally, we note that we must measure that trust *continuously* and not just at the start of a session. For example, actions, such as the act of sending a message, are not limited to the actual process of pressing the button that sends the message, but also include the process of writing the message. *Continuous* or *dynamic* authentication is not a new concept [23], but is a particularly essential one to the paradigm that we propose, which is based on risk and confidence measures that can and often do change over time. Continuous authentication is a critical component to any resilient solution. Such approaches move beyond traditional passwords, crypto-cards, and smart badges, which only provide an simple instant of trust in the current context. Continuous authentication runs in the background authenticating the user regularly (e.g., every keystroke, movement) validating the fact that the user is in the room with a high degree of confidence. This is a simple and very effective part of most face-to-face communications. Continuous authentication may even provide a range of responses, depending on confidence (e.g., not just “allow” or “deny”) [7]. Such a gradated response is an intrinsic and normal part of most face-to-face communications yet rather limited in the digital world. It’s certainly not unreasonable for a bank teller to ask for more identification when a customer’s interactions seem suspicious or they request to move larger sums of money. We believe such continuous and responsive authentication should be an integral part of the digital world.

Simple biometric validation has been exhaustively researched and limitations of such methods are well-known. The continuous, real-time fusion of biometrics to validate identity and a very simple secret to validate intent can be used to improve authentication and reduce ambiguity, improving the probability that the authorized user is who they claim to be and intends to access a resource not just at the start of an authenticated session but throughout it. Simple biometrics are limited in their ability to provide robust authentication. Nevertheless, biometrics are an essential element of any comprehensive authentication scheme.

<sup>2</sup>Auto-destruct: <http://en.memory-alpha.org/wiki/Auto-destruct>

It is important to note that any attempt to provide robust authentication is likely to erode privacy, due to fundamentally contrasting goals. We consider this balance intuitively every time we choose to physically attend a meeting. In addition to verbally sharing our thoughts, physical attendance demands that we reveal much additional information about us. We are for example, revealing what we look and sound like, our location, our mannerisms, and more. In the case of a face-to-face meeting, we nearly unconsciously make the decision that the value of attendance is worth the compromise in privacy.

## 4.2 State and Behavior Metrics

In addition to those metrics that come from a specific, individual biological trait we can add or erode confidence based on combined state and behavioral metrics. Adding combined state and behavioral metrics to a more traditional biometric results in the following notional taxonomy:

1. Something you are
  - (a) Time since certainty\*
  - (b) Spatial-temporal consistency\*
    - i. Time history of biometrics compared to capabilities of the human body.
  - (c) Biometrics (Instantaneous capture/assessment where each of these is a function of each keystroke or other system input)
    - i. Fingerprint typing the letters
    - ii. Keystroke dynamics\*
    - iii. Facial/iris/retina recognition
    - iv. Instantaneous current location (via GPS, camera)
    - v. Voice/grammar/idiom recognition\*
    - vi. Gait
    - vii. Body dynamics (how hold/move my phone)
2. Something you know
  - (a) Password/passphrase
  - (b) Analog combination
3. Something you have
  - (a) Physical key
  - (b) One-time pad
  - (c) RSA SecurID token

Building on the canonical face-to-face interaction we now discuss several such possible combined state and behavioral metrics (marked with a \* in the notional taxonomy above).

**Time since certainty.** The simple fact that time has elapsed since the last face-to-face encounter serves to erode authentication confidence. Some systems may require re-authentication, whether or not the user was active. After a set timeframe model this erosion can be modeled and implemented as a simple step function. In many cases we might be better served by a more gradual erosion of trust.

**Spatial-temporal consistency.** Identity is compromised if a person is perceived to appear in two places at once.

For example, a near-simultaneous appearance in geographically separated locations can indicate that an attack is underway. An individual logging in from to geographically distant locations moments apart could indicate the compromise of a shared secret, such as from a keystroke logger on the first machine that was used to obtain a password to enable login from the second, geographically-distant machine.

Starting from the face-to-face encounter, a spatial probability “basket” can be developed as a function of time based on the equations of state for human mobility. An individual appearing outside this spatial probability basket would erode authentication confidence.

Note that the “consistency” we refer to here is consistency as constrained by the laws of physics and not predictability, for example, if a user is simply trying to authenticate from a place that they do not usually authenticate from, much as is used with credit card fraud detection. Such a measure may be value in some cases but it is anathema to the system that we propose that needs to provide measurable certainty of authentication. Such a system cannot reduce certainty just because of unusual circumstances because the system itself may need to be usable under unusual circumstances.

**Biometrics.** Additionally, a variety of biometrics are frequently not considered in practice as things like fingerprint readers have but also have a measurable and scientific basis:

**Use of grammar and idiom.** Individuals use language differently – speech patterns represent something analogous to a psychological biometric. Grammar checkers analyze text based on accepted rules of proper usage. These grammar rules, however, allow significant freedom. Within these rules, individuals develop specific writing styles that are recognizable. For example, when reading the words “Speak, friend, and enter.” [36], should one interpret this to mean “Say the word ‘friend’ and then enter” or “Friend! Say something and then go inside”? Voice-recognition software exploits this individually specific style to improve speech recognition. In the same way, individual writing style can be used to increase authentication confidence.

**Keystroke dynamics.** Beyond a lingual analysis, keystroke dynamics (and other inputs to computers) have been shown to be a useful tool for authentication [20]. Such a system provides greater confidence with more keystrokes. Enabling such monitoring is yet another real-time continuous metric that can be used to increase or erode trust in a users identity.

As mentioned above, these metrics are fabulously invasive. Thus, they may only be desirable for high assurance, critical systems. Nevertheless, an evaluation similar to that in “The Quest To Replace Passwords” reveals that they would provide dramatic improvements in security with only limited compromises in usability and

deployability.

Our goals combined with the current limitations of science and technology lead to a set of laws that we assert that systems must adhere to in order to measurably capture the elements that make the current gold standard of in-person, human-to-human function well.

## 5. LAWS OF AUTHENTICATION

We posit the *Laws of Authentication* that describe means of measurably validating the amount of trust that one can place in a process of authentication:<sup>3</sup>

0. Identity should be verified as long and as frequently as access to a resource is permitted. If access is ongoing then identity verification should be continuous.
1. Authentication is about measurably validating whether or not someone is who they claim to be, and about determining whether that person intends to authenticate or is being coerced.
2. In person, human-to-human authentication is the “gold standard.” When this is not possible and computers must be involved, then computers should provide measurable measures of confidence (or lack thereof) to humans. Those humans should ultimately make authentication *decisions*, not computers.
3. Authentication should be trivial for the person legitimately authenticating but hard for an adversary to defeat.<sup>4</sup>

Authentication that scales builds on these laws through bootstrapping. Using these techniques, we assert that *humans* should judge things based on the confidence level a computer provides. Moreover, instead of a single sign-in event enabling access until the user logs out, rich authentication intelligently fuses sensor data with predictable human behavior and limitations to enable probabilistic (and difficult to mimic) confidence that the specific user is at the machine. For example, conceivably, every keystroke [19] and mouse motion can be automatically and transparently signed indicating the confidence that the machine input was provided by the specific individual. If confidence is eroded due to user inconsistency (such as injury or sickness), confidence can be restored by requiring more complicated passwords or even more intrusive means such as DNA analysis of blood samples (as such technologies become widely available).

We reiterate that one of the reasons that humans are able to do this is *context*. We observe that in the physical world, *physics* is relevant. Humans can take this into account. Computers can too, but how much should they weigh it? For this reason, we believe authentication should employ additional factors that enable confidence about identity to be increased or decreased. Additional factors could include: *where you are* (e.g., via

<sup>3</sup>With appreciation to Isaac Asimov [3].

<sup>4</sup>Much like verifying a key vs. guessing a key in public key cryptography.

GPS) [12]—because human motion is governed by the laws of physics; and *time*—it is physically impossible for a human to be in two places at once, and so if Person A was in California at 10:00 AM on Tuesday, they won’t be in Beijing one hour later.

Moreover, we observe that in the physical world, “authentication” (recognizing someone) draws heavily on intuition. In fact, humans perform this function naturally and trivially. In contrast, machines cannot understand the subtlety of authentication. For example, they cannot easily interpret the meaning of being “Facebook friends” with someone. Again, in contrast, humans can recognize gait, posture, and other forms of “body language,” even via videoconferencing.

Alternatively, suppose that there exists background noise in a phone call from Disneyland. Humans recognize such sounds and intuitively and unconsciously check to ensure that this background information is consistent with the authentication “picture.” For example, if an individual in a phone call claims to be delayed at the airport, sounds that are inconsistent with an airport (i.e., sounds from Disneyland) would serve to erode confidence in this authentication picture. Other participants in the phone call may choose to increase confidence in identity by pointing out this inconsistency. A simple and unobtrusive inquiry (“If you are at the airport, why do I hear Disneyland sounds in the background?”) followed by a credible response (“Oh, I’m walking past the Disney store in the airport lobby right now.”) would serve to increase confidence. Therefore, our idea is to present rich information to other human users to enable detection of inconsistencies in authentication.

Note that the “consistency check” provided by such background noise is anathema to computing systems. For example, Google Voice transcribes phone messages into text which is then e-mailed to the recipient. Background noise in a phone message complicates the speech recognition process and is unnecessary for the transcription to text. Therefore, the computing system regards such background noise as a complicating factor to be filtered out. As a result, the transcribed message loses a rich source of authentication information that could provide a consistency check.

## 6. EXAMPLES

In this section, we consider two alternative applications for continuous or dynamic authentication.

**Human-Machine Authentication.** Continuous authentication develops a probability of user presence that moves up and down depending on the situation and as new sources of data are added. As a user walks into the same room with the “authenticator device” (a computer, tablet, smart phone, etc.), the authenticator device may use a facial recognition algorithm to develop a probability of the user being in the room of 0.7. As the user picks up the authenticator device, the device may use

accelerometers and rate gyros to recognize gestures invoked by the user to increase this probability to 0.85. As the user moves around, the authentication device may use the same devices to increase the probability that the user is in the room to 0.93 over time. After two hours, continued observation of the user’s biometrics may increase the probability that the user is in the room to 0.97. When the user puts on the iris/retina recognition/eye-tracking glasses, the probability of the user’s presence goes to 0.993. During all of this, the user might be dictating to a computer. Based on the fact that the human does not use any “duress words” during that time, the probability of duress is low. After the user sends a couple of emails, observation of the user’s typing rhythms provides a probability that those emails come from the user and that the user is sanguine goes to 0.99995.

**Human-Human Authentication.** Humans are far more capable authenticators than machines. Nevertheless, the assessment described above can provide useful information to aid human authentication. For many interactions, the authentication assessment prepared for human-machine authentication may be more than sufficient. In addition, the human-machine authentication can keep the user from making stupid mistakes.

For critical interactions in a high-threat environments, however, raw information (such as the video feed of the user walking into the room, etc.) should be provided with high integrity to the decision-maker. In such situations, inconsistencies between the human-machine authentication result and the raw information provided to the decision-maker may serve to erode confidence below that developed through human-machine authentication. In the same way a heads up display helps pilots track aircraft status and select targets we envision an environment where a machine may help the humans process the data but in the end they select the target and they verify the authentication.

**Machine-Human Authentication.** We note in passing that there are examples of where a machine may need to authenticate to a human, for example a remote sensor in a hostile environment. While this is a related issue to what we discuss in this paper because it also involves *origin integrity*, the topic is mostly out of the scope of this paper because we believe that the solutions relate less to measurably verifying identity and more to the integrity of the data and/or sensor. Thus, we feel machine-human authentication relates more to a combination of data provenance, our previously-described approach on clean slate designs [30], and, in some cases, procedures [28] similar to zero-knowledge protocols [14].

## 7. SUMMARY

Online activities can approach the level of clarity, certainty and intuitiveness as activities in the physical world. Physical world metaphors drive the entire user expe-

rience. However, the misapplication of some of these metaphors physical metaphors—e.g., resemblance as opposed to mere consistency—can create anxiety and a lack of clarity for users about online actions. Moreover, a mismatch in goals—e.g., preventing attack of a captured set of password hashes, rather than validating user identity—lead to solutions that are inappropriate in some situations, and certainly in critical environments. Our laws of authentication are a solution to this mismatch. By properly ensuring consistency between two worlds and appropriately managing the role of humans vs. the role of computers, the “membrane” between the physical and online world effectively disappears.

## Acknowledgements

This research was supported in part by the Director, Office of Computational and Technology Research, Division of Mathematical, Information, and Computational Sciences of the U.S. Department of Energy, under contract number DE-AC02-05CH11231. It is also supported in part by the National Science Foundation and the GENI Project Office under Grant Number CNS-0940805, and by the National Science Foundation under Grant Numbers CCF-1018871 and CCF-1049738.

Some of this work was completed while Ed Talbot was a Distinguished Member of Technical Staff at Sandia National Laboratories. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000.

Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect those of any of the employers or sponsors of this work.

## 8. REFERENCES

- [1] P. S. Aleksic and A. K. Katsaggelos. Audio-Visual Biometrics. *Proceedings of the IEEE*, 94(11):2025–2044, Nov. 2006.
- [2] M. Ambinder. Why Clinton’s Losing the Nuclear Biscuit Was Really, Really Bad. *The Atlantic*, Oct. 22 2010.
- [3] I. Asimov. Runaround. *Astounding Science Fiction*, March 1942.
- [4] BBC News. Clinton drops nuclear football. <http://news.bbc.co.uk/2/hi/americas/328442.stm>, Apr. 26 1999.
- [5] S. Bellovin and G. McGraw. Silver Bullet Show 081 - An Interview with Steve Bellovin. <http://www.cigital.com/silver-bullet/show-081/>, Dec. 26 2012.
- [6] M. Bishop. *Computer Security: Art and Science*. Addison-Wesley Professional, Boston, MA, 2003.
- [7] M. Bishop, S. Engle, D. A. Frincke, C. Gates, F. L. Greitzer, S. Peisert, and S. Whalen. A Risk Management Approach to the ‘Insider Threat’. In C. W. Probst, J. Hunker, and M. Bishop, editors,

*Insider Threats in Cybersecurity*, Advances in Information Security Series, pages 115–138. Springer, Berlin, September 2010.

[8] M. Bishop, S. Engle, S. Peisert, S. Whalen, and C. Gates. We Have Met the Enemy and He is Us. In *Proceedings of the 2008 New Security Paradigms Workshop (NSPW)*, Lake Tahoe, CA, September 22–25, 2008.

[9] J. Bonneau, C. Herley, P. C. van Oorschot, and F. Stajano. The Quest to Replace Passwords: A Framework for Comparative Evaluation of Web Authentication Schemes. In *Proceedings of the 33rd IEEE Symposium on Security and Privacy*, pages 553–567, May 20–23, 2012.

[10] D. Chaum. Security Without Identification: Transaction Systems to Make Big Brother Obsolete. *Communications of the ACM*, 28(10):1030–1044, 1985.

[11] J. Daugman. How Iris Recognition Works. *IEEE Transactions on Circuits and Systems for Video Technology*, 14(1):21–30, Jan. 2004.

[12] D. E. Denning and P. F. MacDoran. Location-Based Authentication: Grounding Cyberspace for Better Security. *Computer Fraud & Security*, 1996(2):12–16, 1996.

[13] S. Egelman, D. Molnar, N. Christin, A. Acquisti, C. Herley, and S. Krishnamurthi. Please Continue to Hold: An Empirical Study on User Tolerance of Security Delays. In *Proceedings of the Ninth Workshop on the Economics of Information Security (WEIS)*, Harvard University, Cambridge, MA, June 7–8 2010.

[14] S. Goldwasser, S. Micali, and C. Rackoff. The Knowledge Complexity of Interactive Proof-Systems. In *Proc. of the 17th Annual ACM Symposium on Theory of Computing*, 1985.

[15] C. Herley. So Long, and No Thanks for the Externalities: The Rational Rejection of Security Advice by Users. In *Proc. of the 2009 New Security Paradigms Workshop (NSPW)*, 2009.

[16] G. Holland and L. Burton et al. Dreadnought (Star Trek: Voyager). Paramount, February 1996.

[17] F. Hsu, H. Chen, and S. Machiraju. WebCallerID: Leveraging Cellular Networks for Web Authentication. *Journal of Computer Security*, 19(5):869–893, 2011.

[18] A. K. Jain, A. Ross, and S. Pankanti. Biometrics: a Tool for Information Security. *IEEE Transactions on Information Forensics and Security*, 1(2):125–143, 2006.

[19] K. S. Killourhy and R. A. Maxion. Comparing Anomaly-Detection Algorithms for Keystroke Dynamics. In *IEEE/IFIP International Conference on Dependable Systems & Networks (DSN)*, pages 125–134, 2009.

[20] K. S. Killourhy and R. A. Maxion. Free vs. Transcribed Text for Keystroke-Dynamics Evaluations. In *Proceedings of the 2012 Workshop on Learning from Authoritative Security Experiment Results*, LASER '12, pages 1–8, New York, NY, USA, 2012. ACM.

[21] M. Kumar, T. Garfinkel, D. Boneh, and T. Winograd. Reducing Shoulder-Surfing by Using Gaze-Based Password Entry. In *Proceedings of the 3rd Symposium on Usable Privacy and Security (SOUPS)*, pages 13–19, 2007.

[22] A. J. Lee and M. Winslett. Enforcing Safety and Consistency Constraints in Policy-Based Authorization Systems. *ACM Transactions on Information and System Security (TISSEC)*, 12(2):1–33, 2008.

[23] J. Leggett, G. Williams, M. Usnick, and M. Longnecker. Dynamic Identity Verification via Keystroke Characteristics. *International Journal of Man-Machine Studies*, 35(6):859–870, 1991.

[24] T. F. Lunt and R. Jagannathan. A Prototype Real-Time Intrusion-Detection Expert System (IDES). In *Proceedings of the 1988 IEEE Symposium on Security and Privacy*, pages 59–66, Oakland, CA, April 18–21, 1988.

[25] A. Mehrabian. *Silent Messages*. Wadsworth, 1971.

[26] A. Narayanan, H. Paskov, N. Z. Gong, J. Bethencourt, E. C. Emil Stefanov, R. Shin, and D. Song. On the Feasibility of Internet-Scale Author Identification. In *Proceedings of the 33rd IEEE Symposium on Security and Privacy*, San Francisco, CA, May 20–23, 2012.

[27] Y. Niu and H. Chen. Gesture Authentication with Touch Input for Mobile Devices. In *Proceedings of the 3rd International Conference on Security and Privacy in Mobile Information and Communication Systems (MobiSec)*, May 2011.

[28] Office of the Director of National Intelligence. United States Intelligence Community Information Sharing Strategy. [http://www.dni.gov/reports/IC\\_Information\\_Sharing\\_Strategy.pdf](http://www.dni.gov/reports/IC_Information_Sharing_Strategy.pdf), Feb. 22 2008.

[29] L. O’Gorman. Comparing Passwords, Tokens, and Biometrics for User Authentication. *Proceedings of the IEEE*, 91(12):2021–2040, Dec. 2003.

[30] S. Peisert, E. Talbot, and M. Bishop. Turtles All The Way Down: A Clean-Slate, Ground-Up, First-Principles Approach to Secure Systems. In *Proceedings of the 2012 New Security Paradigms Workshop (NSPW)*, pages 15–26, Bertinoro, Italy, September 19–21, 2012.

[31] G. Roddenberry and G. L. Coon et al. Bread and Circuses (Star Trek: The Original Series). NBC, March 1968.

[32] C. Song, Z. Qu, N. Blumm, and A.-L. Barabási. Limits of Predictability in Human Mobility. *Science*, 327(5968):1018–1021, 2010.

[33] S. Spielberg and G. Lucas, et al. Indiana Jones and the Last Crusade. Lucasfilm Ltd., 1989.

[34] S. D. Spray and J. A. Cooper. The Unique Signal Concept for Detonation Safety in Nuclear Weapons. Technical Report SAND91-1269, Sandia National Labs., Albuquerque, NM (United States), June 1993.

[35] E. B. Talbot, D. Frincke, and M. Bishop.

Demythifying Cybersecurity. *IEEE Security and Privacy*, 8(3):56–59, May/June 2010.

[36] J. R. R. Tolkien. *The Fellowship of the Ring*. George Allen & Unwin, 1954.

[37] A. M. Turing. Computing Machinery and Intelligence. *Mind*, 59(236):433–460, 1950.