

Open Access License Notice

This article is © its author(s) and is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). This license applies regardless of any copyright or pricing statements appearing later in this PDF. Those statements reflect formatting from the print edition and do not represent the current open access licensing policy.

License details: <https://creativecommons.org/licenses/by/4.0/>

A Perceptual Taxonomy of Contextual Cues for Cyber Trust

Abstract

People are the weakest link in the security of any computer system or network. Adversaries tactically exploit decision-making processes of computer users by preying on their propensity to trust electronic communications based on learned contextual cues and environmental influences. Despite this understanding, explorations into the psychological, experiential, technological, and environmental factors influencing the cognitive phenomena of trust and suspicion are not typically based on the cues available to computer users. This work defines a taxonomy of human-perceptible trust cues that are used to make decisions online across the domains of Web, E-mail, and Social Networking. We believe this taxonomy could be foundational for future studies of phishing attacks. The Taxonomy also has significant implications for training both cyber security personnel and general computer users, because it codifies the domain or body of knowledge that may become the basis for future development of information systems security education and training.

1. Introduction

The most mercurial element of any information system is the human component. Security professionals and hackers have long known that end-users are the weakest link in any organization's security posture [14]. Regardless of resources invested in technological controls to mitigate cyber-attacks, organizational security ultimately rests in the hands of the end-user [10].

Cyber attacks have been classified as *physical, syntactic, or semantic* [5]. Phishing attacks are semantic attacks, in which victims are conned by fraudulent e-mails, web pages, and social network accounts. Such attacks have become widespread with the advent of ubiquitous computing and pose a serious threat to users, and to organizations that lose public trust as a consequence [16]. Victims perceive that phishing attempts, generally emanating from one of three modalities -- e-mail, social networks, and the web -- originate from trusted sources, when in reality they come from con artists [5]. Attackers capture passwords, credit card numbers, and other personal information from their victims.

Technical counter measures to these attacks are only partially effective. Much of the responsibility for detecting phishing and social engineering attacks rests with end-users [16]. Security is strengthened by well-designed technology and systems, but also through user training [2]. Effective strategies for guarding against cyber threats combine technological and human detectors [16].

Alert to cues and lures, experienced users are better equipped to discern whether a message is trustworthy. Less experienced users cannot identify lures and cues associated with untrustworthy e-mails or other computer-mediated messages [17]. Unfortunately, current methods of security awareness training have proven less than effective. Training is often too broad, perfunctory, and sporadic to develop the competencies required to operate safely on a network. A company is only as secure as its weakest link, which is all too often an untrained workforce. In order to effectively train online users to recognize observable phishing cues and lures that commonly appear in websites, e-mails, and social networks, we must first understand the *trust surface* across these modalities. This study describes a human-perceptible *Cyber Trust Taxonomy* to classify modality specific trust elements commonly leveraged by attackers.

2. Background

Social engineering attacks soliciting user information or enticing users to carry out risky actions online are known as phishing attacks [4]. Such attempts are carried out in various ways across three principal online domains -- e-mail, web, and social networks. Generally, these attacks attempt to look, feel, and sound legitimate to the unsuspecting user. In some cases the legitimacy of a phishing attempt may be visually indistinguishable from a legitimate form of online communication [4]. This makes effective online security awareness training important.

Behavioral measures meant to predict an individual's likelihood to trust are often grouped into two categories: experience and disposition. Disposition refers to an individual's inherent characteristics which are determined by numerous biological and environmental factors over the course of years. These characteristics include things like propensity to trust, honesty, and risk aversion. In contrast to dispositional characteristics, which broadly apply to different areas, experience is domain specific [15]. For example, expertise with computers is likely to affect an individual's trust in the cyber domain, but it is unlikely to affect trust in face-to-face communication. In one study, 299 subjects were given questionnaires measuring experiential factors including computer self-efficacy, web experience, and security knowledge as well as dispositional factors including trust, risk aversion, and suspicion [18]. Each participant was a student in an information systems class and was given a unique code used to access course materials and tests. They completed coursework covering internet security and phishing and were reminded daily never to divulge their codes to anyone. At the end of the course, a phishing email was sent to each subject requesting their secret code. Approximately one third of the subjects failed to recognize the phishing attack and responded with their codes. The main result from this study was that experiential factors had a much larger impact on phishing success than dispositional ones, indicating the impact training might have on improving cyber security.

Research on phishing explores the causes of susceptibility to phishing attacks. [16] describes an integrated information processing model for explaining individual differences in phishing attack susceptibility. The model considers variations in susceptibility as illuminated by habitual use of online media, cognitive load, and other aspects of user activity. Investigations on the relative impact of awareness and security indicator controls to mitigate phishing hazards have at once demonstrated the need for more effective controls and revealed the limitations of our understanding of the interplay between users and information technology with respect to trust decisions [7] [11] [19].

In order to classify artifacts of trust that compose phishing content, we propose a taxonomy focused on cyber trust. A taxonomy provides a method to unambiguously classify related items into distinct domains. A taxonomy is necessary in order to create a common vocabulary and an understanding of trust cues online [12]. This allows researchers to study individual trust cues that make up a larger whole. To our knowledge, a classification scheme such as this has yet to be applied to perceivable elements of trust within the three online domains of web, email, and social media.

3. Taxonomy

In an effort to better understand phishing from the user perspective, we have developed a taxonomy of lures and cues. The taxonomy serves as a framework for research, online interventions, and user training. The first step to understanding why users fall victim to phishing attacks is to classify the individual components that make up these attacks. This work focuses on the human-observable elements of phishing attacks. By categorizing trust artifacts within a phishing message, we can begin to study the visual elements used by attackers to deceive end-users. We approach this work from the perspective of how attackers exploit end-user trust propensity. In general, trust decisions users make online are based on three distinct hazard cue categories: *Content*, *Context*, and *Contract*; (i) Trust decisions based on visual **content** presented to the subject, (ii) Trust decisions based on the **context** in which content elements are presented to the subject, and (iii) An implied trust **contract** is made between the victim and the attacker when a victim is asked to supply sensitive information or to perform a risky action by the attacker.

These hazard cue categories provide a framework with which we can further analyze and classify phishing artifacts across the three aforementioned online modalities. A discussion of the three categories of cues follows. Figure 3-1 presents an overview of the perceptual cyber trust taxonomy.




Figure 3-1: Cyber Trust Taxonomy Hierarchy

3.1 Content

Visual artifacts that make up the phishing content with which users are confronted are the front line tools that attackers employ to manipulate user trust. Many users are trained to look for visual security cues, such as HTTPS or padlocks, within a web browser to verify the presence of "security" online. Unfortunately, attackers leverage such reliance on the presence of expected visual indicators, and exploit this to their advantage by embedding such elements in web pages or pop-ups using visual deception. Attackers have also been known to take advantage of the user's lack of knowledge about the syntactical structures of URL domain names, filenames in email attachments, and the reliance on reputable names or logos. Figure 3-2 illustrates the *Content* branch of our perceptual taxonomy.

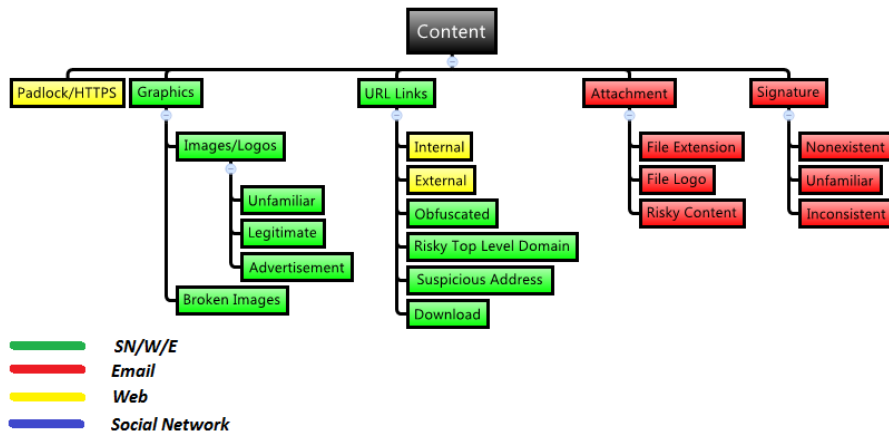


Figure 3-2: Content Artifacts of Trust

Padlock/HTTPS: These are binary elements within a web browser that indicate two things—a secure connection in the form of an encrypted connection between the user’s browser and a web server, and the endorsement of the associated session key on behalf of a certificate authority (CA) [9]. Web browsers, such as Google Chrome, indicate to the user whether or not the CA used for end-to-end encryption is trusted and reputable (Figure 3-3).

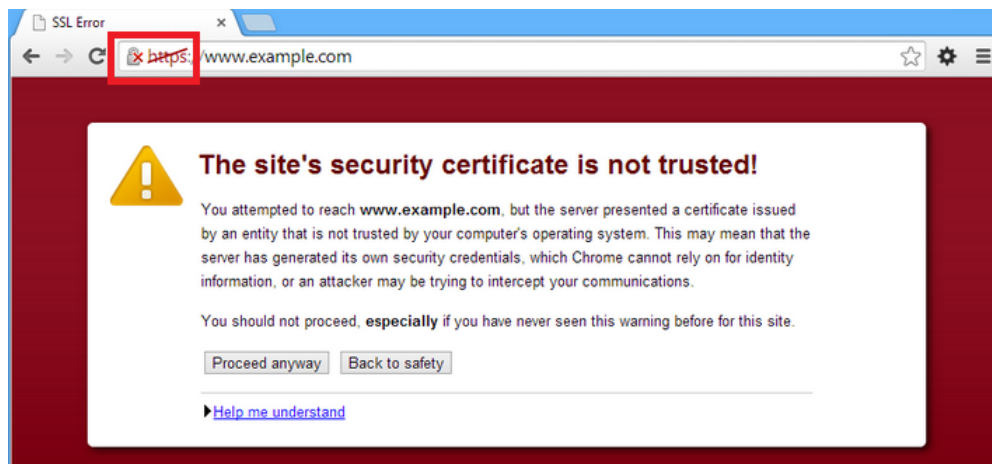


Figure 3-3: Google Chrome Alerting the User of an Untrusted SSL Certificate

Unfortunately, there is no universally agreed upon method for how these visual indicators should be presented to the user. They are consequently presented in various ways depending on browser vendor and version [7] [10]. Attackers can tailor visual padlock deceptions to mimic those of a user's browser [1] [9]. Moreover, the variability in cue forms and presentations add to the general confusion of users.

Graphics: The presence or absence of expected images online can influence user trust propensity. Subtle discrepancies in expected logos and images on web pages can have direct influence on user trust propensity. Alternatively, the use of legitimate logos within a phishing site can increase user trust propensity. The perceived production quality or lack thereof in logos and images can also represent cues to users as to the trustworthiness of online content.

URL Links: Universal Resource Locators (URLs) have syntactical cues that can be evaluated. The ability to inspect a URL and readily recognize whether it links to a trusted web site is vital, as URL misdirection is a common tactic for steering users to compromised web pages [13]. URLs that have been obfuscated to hide their real target page are also used to mislead naive users [3]. Legitimate URL aliases, such as TinyURL may confound cue recognition. Evaluation of top level domains (TLDs) at the end of the URL should be carried out before clicking on a link. Domains ending with a risky TLD name such as .RU, .CH, and .KP should be evaluated as extremely risky by the end-user [13]. Finally, a URL can be inspected to gather a general idea as to its action if the user clicks on it. The URL could be a link to immediately download a file. In this case, the file type provides an additional trust/hazard cue.

Attachments: The logos and file extensions of email attachments can be evaluated to identify hazard cues. Risky content in the form of an email attachment should be considered as any executable file or multi-media file that is processed within a local client application. Visual trickery with the way certain email clients render file names can be used to deceive users into thinking they are opening up a PDF file, when in actuality the file is an executable (.exe) file. The file attachment logo can also be used as an observable cue for file type identification.

Email Signature: The mere presence or absence of an email signature block can impact whether recipients trust an e-mail. Subtle discrepancies between signature blocks and other information, such as e-mail header content, can also cause a user to be wary of the legitimacy of an email message.

The presence or absence of visual content trust cues online can provide some level of reassurance to the user about the trustworthiness of an online engagement. These binary indicators by themselves do not necessarily guarantee a secure state in which an application can be completely trusted, but rather a level of trust that when used within the context of other elements can provide more insight into the trustworthiness of an online engagement.

3.2 Context

The context in which phishing artifacts are presented can affect user trust propensity as well. The presentation of visual content can influence the focus of users based on designer end goals. Context combines the visual presentations of content with user expectations arising from environmental and historical factors. The origin of an e-mail, URL, or follower can have significant trust implications. The tone of an e-mail or social media message can also affect the user's willingness to trust. This is commonly dependent upon the user's relationship to the other party in question. Grammatical, syntactical and idiomatic characteristics of sentences, words and phrases across all three online modalities may affect user trust. Figure 3-4 presents the *Context* branch of the Cyber Trust taxonomy.

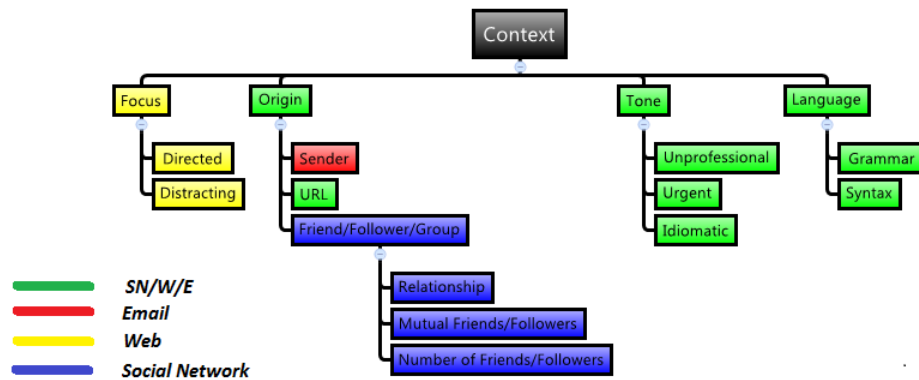


Figure 3-4: Context Artifacts of Trust

Focus: Positioning of pictures and text on a web page can impact user focus and consequently influence trust propensity. A user trying to accomplish a given task (e.g., downloading a file from a download page) can become easily distracted by web pages containing multiple images and links that read "download" and are cluttered with advertisements. Other web sites may have a more intuitive flow allowing a user to be directed to easily accomplish the task at hand with little if any distractions.

Origin: The domain name of the sender address can be used with other trust elements to deduce the trustworthiness of an e-mail message [6]. The same is true of the TLD of a URL. Information related to the relationship to the user, the number of friends/followers, and mutual friends/followers in a social network can be used to derive the legitimacy of an individual or organization via social media networks.

Tone: Tone in the contextual sense refers to the way in which users are addressed via online communication. For example, one might expect a business associate to begin an e-mail with a more formal greeting than a close friend. Conversely, one might not expect a friend to exhibit formality throughout a message. A sense of urgency in a message is often employed by attackers to compel naive users to supply sensitive data or carry out a risky action, such as opening a malicious attachment.

Language: Users expect senders and recipients to exhibit some level of predictability with regard to grammar, vocabulary, and sentence structure. Deviations from anticipated language usage may influence user trust propensity in either a positive or negative way.

3.3 Contract

This portion of the taxonomy is concerned with the value proposition of an online trust decision. A *Contract* is composed of an offer and consideration. Consideration captures what users are giving in exchange for what they are getting in return (offer). Users are commonly asked to part with sensitive information or execute some action that may place them in some heightened state of vulnerability. The

offer ostensibly confers the benefit of complying with the consideration component of the online contract. Figure 3-5 presents this branch of the taxonomy.

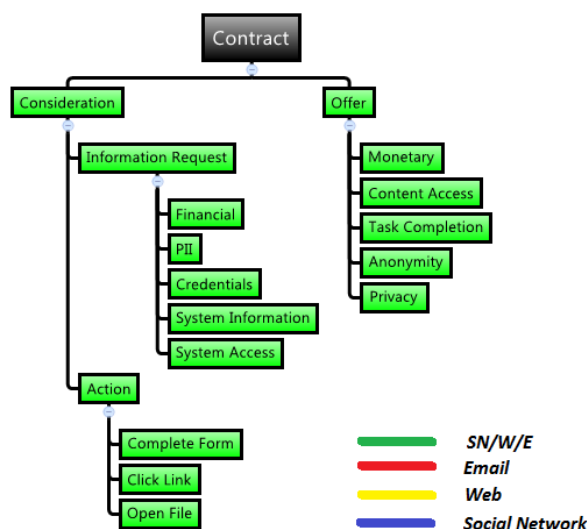


Figure 3-5: Contract Artifacts of Trust

Consideration: Attackers entice users to provide Personal Identifiable Information (PII), credentials, system access, or general system information in order to further their attacks. Alternatively, an attacker will try to persuade the victim into performing a risky action that can lead to disclosure of sensitive information or system compromise. System compromise can occur when a user is enticed to navigate to a compromised web site, or when a user opens a malicious file.

Offer: This branch is focused on classifying what the attacker is advertising to provide the user in exchange for information. Attackers may entice users with large sums of money, sensitive information, or access to subscription-based content or services.

4. Applications to Research, Tool Development and Training

Our perceptual taxonomy of cyber trust cues embodies a framework within which to pose new research questions, develop novel tools, and pursue differentiated training for individuals and organizations. Common to these activities is the use of the taxonomy to classify, sort, and otherwise analyze phishing attacks by its constituent and composed elements. We believe this taxonomy will ultimately enhance information system security education. This section describes opportunities for applying the taxonomy to each domain and establishes corresponding pathways of their pursuit.

4.1 Classification

With respect to research in this field, the taxonomy helps classify phishing attacks pulled from the wild. Against the backdrop of a stable collection of phishing attack categories, several research questions can be posed. *Which attacks are the most effective, and against what targets? What kinds of attack are amenable to mitigation by technical controls? Which are most effectively dealt with by targeted training?*

Classification may be conducted on an *existential* basis, in which artifacts are inspected for the inclusion of specific elements of the taxonomy. Investigators may extend this approach by establishing *types* or values for individual elements for more meaningful comparisons and matching. *Combinational* classification schemes can also be used that, for instance, group attacks sharing an urgent tone with one that asks a user to provide credentials in exchange for the promise of successful completion of a task.

Classifying individual phishing artifacts onto the cyber trust taxonomy provides a basis for clustering similar artifacts. Taxonomical elements in phishing artifacts permit “fingerprinting” – capturing the distinctive aspects of an attack. Clustering artifacts with similar or identical fingerprints may offer some insights as to the origin or development process of an attack.

Classification can also be used to profile the phishing attacks observed by an individual or organization. This yields an added level of intelligence concerning the threat landscape of an enterprise that can be used to optimize risk-reducing activities. Such classifications can help a security team develop tailored responses and interventions that target critical facets of the characteristics common to observed attacks. For example, based on the prevalence of phishing attacks targeting personally identifiable information (PII) procedural or operational controls may be deployed to gather PII for legitimate purposes using a secure/specialized channel, obviating the need to respond to emails requesting PII.

Developing sound and useful classification techniques and schemes for phishing must be a research priority to advance the core science and understanding of the topic. Any such agenda should leverage existing phishing artifact repositories, classifying and analyzing the classified specimens therein. Analysis must be conducted against a core set of measurable attributes or effects of phishing attacks to address the major research questions articulated above. When it comes to understanding the human factors involved, this represents an opportunity to define experiments designed to study the influence of psychology and neurobiology on cyber trust as directed against distinct taxonomical elements.

4.2 Tool Development

Our taxonomy establishes a blueprint for the development of programmatic tools to facilitate targeted interventions for end-user trust decisions. Assisting in the development of such tools is an XML document type definition (DTD) based on the taxonomy. Figure 4-1 provides an example showing the content attributes of phishing email that contains an external URL link to a suspicious site embedded within the message. This example also includes the attachment of a risky PDF document.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE CYBER_TRUST_TAXONOMY SYSTEM "CyberTrust.dtd">
- <CYBER_TRUST_TAXONOMY>
  - <CONTENT>
    <PADLOCK>Yes</PADLOCK>
    - <GRAPHICS>
      - <IMAGES_LOGOS>
        <UNFAMILIAR>No</UNFAMILIAR>
        <LEGITIMATE>Yes</LEGITIMATE>
        <ADVERTISEMENT>Yes</ADVERTISEMENT>
      </IMAGES_LOGOS>
      <BROKEN_IMAGES>No</BROKEN_IMAGES>
    </GRAPHICS>
    - <URL_LINKS>
      <INTERNAL>No</INTERNAL>
      <EXTERNAL>Yes</EXTERNAL>
      <OBFUSCATED>No</OBFUSCATED>
      <RISKY_TLD>Yes</RISKY_TLD>
      <SUSPICIOUS_ADDRESS>Yes</SUSPICIOUS_ADDRESS>
      <DOWNLOAD>Yes</DOWNLOAD>
    </URL_LINKS>
    - <ATTACHMENT>
      <FILE_EXTENSION>Pdf</FILE_EXTENSION>
      <FILE_LOGO>Pdf</FILE_LOGO>
      <RISKY_CONTENT>Yes</RISKY_CONTENT>
    </ATTACHMENT>
  </CONTENT>
+ <CONTEXT>
+ <CONTRACT>
</CYBER_TRUST_TAXONOMY>
```

Figure 4-1: XML Describing Phishing Email

One such example is a tool to automatically generate and render phishing content across online modalities. This tool produces static content in the form of .HTML files, given the quantity of stimuli to produce for each modality (Figure 4-2). The tool relies on cascading style sheets (.CSS) that define the layout of individual content artifacts that define a web page, e-mail, or social network page (Figure 4-2). This tool also relies on the presence of a pool of textual messages and images with which to generate and

degrade content. These artifacts are stored in a local SQLite3 database. The cyber trust content generator produces content with random levels of artifact degradation based on our taxonomy.

Future plans for tool development using the taxonomy should complement the foundational research agenda. Of initial consideration are tools for measuring attributes of phishing (and legitimate) specimens, thus the definition of meaningful metrics becomes an immediate objective. Tools that permit measurement and exploration of such content will facilitate research and ultimately translate to superior detective controls and technical interventions within browsers, applications and operating systems.

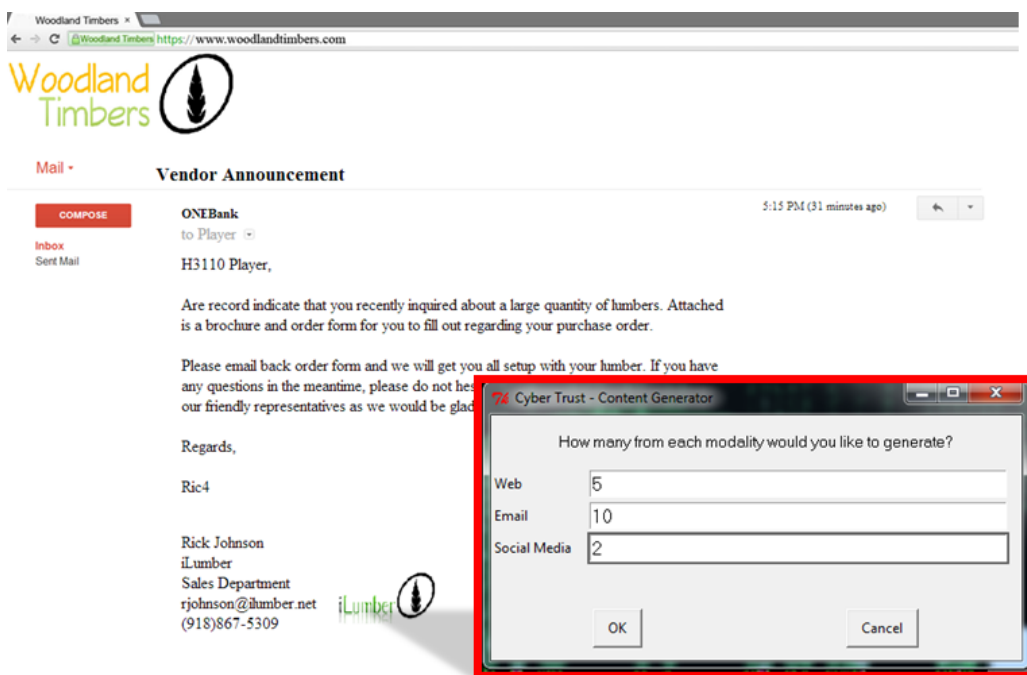


Figure 4-2: Email Stimulus Produced by Content Generator

4.3 Targeted Training

In addition to driving research and promoting tool construction, this taxonomy is useful for developing differentiated training program content designed to strengthen end-user cyber security competencies. The goal is to enable users to detect deceitful messages, dodge their deleterious effects, and disseminate information about deceit to stakeholders. But, how is the goal disaggregated into its trainable facets? The first step involves conducting a learning needs assessment. Training programs are designed to achieve goals that meet instructional needs. It is dangerous and costly to begin any program without a complete assessment of the task, behaviors, and environment [8]. Understanding workforce capabilities is a critical part of identifying areas requiring change. A learning needs assessment involves asking questions that reveal the competencies and development needs of end-users. The taxonomy provides a domain from which to draw these questions. One-size-fits-all training is ill-advised, given the broad continuum of end-user skill-levels and dispositions. Knowledge of end-user strengths and development needs allows organizations to implement differentiated training strategies designed to equip end-users with the requisite know-how to evade cyber-deception. It is a vital step toward creating a frontline defense for otherwise vulnerable organizations.

The taxonomy not only guides the development of end-user learning content, but also provides a blueprint for training program evaluation. The taxonomy establishes criteria for training success. It

provides organizations with an organized means of determining whether employees have overcome the learning deficiencies identified by the needs assessment.

To illustrate a training application, consider the taxonomical node URL Obfuscation. To address an organizational concern over phishing, a company administers a taxonomically-oriented test to its employees. The results indicate that employees struggle distinguishing trustworthy URLs from obfuscated URLs. They are vulnerable to typejacking domain name attacks. This type of deception involves surreptitiously altering a legitimate domain name (e.g., www.paypal.com vs. www.paypal.com).

Having determined that a company is vulnerable to URL Obfuscation, the next step is to establish a remedial training objective and module addressing this type of trick and the adverse consequences of overlooking it. This training module is then applied to those who have been identified as deficient in URL Obfuscation detection. Applying a differentiated training strategy is important, because there is no need to train employees who are already competent in detecting URL Obfuscation.

To evaluate the effectiveness of URL Obfuscation training, one can use a parallel form of the employee test administered in the needs assessment. By sending out periodic decoy messages with obfuscated URLs, an organization may determine whether employees are applying what they have learned. In addition, an organization can evaluate the return on investment by examining the reduction in costs associated with URL Obfuscation detection failure.

Beyond the taxonomy itself, it is vital to assess whether the learning process has inspired users to continue their learning. Did that which was learned transfer to the jobs they perform? What observable effects resulted from training (e.g., fewer cyber-attack breaches, fewer cyber-related monetary and customer losses)? Evaluative information from a well-planned training program provides valuable feedback for continuous improvement of training program content, methods, outcomes, and results.

The key to unleashing the potential of the taxonomy in this application domain is building instructional modules tuned to training against specific cues and lures identified within it. As a adjunct with previously stated research objectives, such training modules can be evaluated within the context of human subject studies to evaluate their efficacy in targeted applications. -This must be done in coordination with an identified learning model and framework that offers students and users the best opportunity to absorb and retain the content.

5 Conclusions

This paper describes the development and structure of a taxonomy for human-observable cues for making trust decisions online. This taxonomy has the potential to facilitate research, aid in developing security tools, and enhance the training of end-users within organizations. Future work will focus on the development of simulated phishing content generation across all modalities. With the assistance of an eye tracker, we will test both technically naive and savvy subjects with the generated stimuli to identify which trust cues are most successful for manipulating end-user trust in cyberspace.

Acknowledgements

This material is based on research sponsored in part by the Air Force Office of Scientific Research (AFOSR), under award number FA9550-12-1-0457. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the AFOSR.

References

1. Adelsbach, A., Gajek, S., & Schwenk, J. (2005). Visual spoofing of SSL protected web sites and effective countermeasures. In *Information Security Practice and Experience* (pp. 204-215). Springer Berlin Heidelberg.

2. Bowen, B. M., Devarajan, R., & Stolfo, S. (2011, November). Measuring the human factor of cyber security. In *Technologies for Homeland Security (HST), 2011 IEEE International Conference on* (pp. 230-235). IEEE.
3. Cova, M., Kruegel, C., & Vigna, G. (2008). There is no free phish: an analysis of "free" and live phishing kits. In *Proceedings of the 2nd conference on USENIX Workshop on offensive technologies (WOOT'08)*. USENIX Association, Berkeley, CA, USA, , Article 4 , 8 pages.
4. Dhamija, R., Tygar, J. D., & Hearst, M. (2006, April). Why phishing works. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 581-590). ACM.
5. Downs, J. S., Holbrook, M. B., & Cranor, L. F. (2006, July). Decision strategies and susceptibility to phishing. In *Proceedings of the second symposium on Usable privacy and security* (pp. 79-90). ACM.
6. Drake, C. E., Oliver, J. J., & Koontz, E. J. (2004, July). Anatomy of a Phishing Email. In *CEAS*.
7. Egelman, S., Cranor, L. F., & Hong, J. (2008, April). You've been warned: an empirical study of the effectiveness of web browser phishing warnings. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1065-1074). ACM.
8. Goldstein, I. L., & Ford, J. K. (2011). *Training in Organizations*. Belmont, CA: Wadsworth.
9. Herzberg, A., & Jbara, A. (2008). Security and identification indicators for browsers against spoofing and phishing attacks. *ACM Transactions on Internet Technology (TOIT)*, 8(4), 15.
10. Hong, J. (2012). The state of phishing attacks. *Communications of the ACM*, 55(1), 74-81.
11. Kumaraguru, P., Cranshaw, J., Acquisti, A., Cranor, L., Hong, J., Blair, M. A., & Pham, T. (2009, July). School of phish: a real-world evaluation of anti-phishing training. In *Proceedings of the 5th Symposium on Usable Privacy and Security* (p. 3). ACM.
12. Landwehr, C. E., Bull, A. R., McDermott, J. P., & Choi, W. S. (1994). A taxonomy of computer program security flaws. *ACM Computing Surveys (CSUR)*, 26(3), 211-254.
13. McGrath, D. K., & Gupta, M. (2008). Behind phishing: an examination of phisher modi operandi. In *Proceedings of the 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats (LEET'08)*, Fabian Monrose (Ed.). USENIX Association, Berkeley, CA, USA, , Article 4.
14. Ryan West. 2008. The psychology of security. *Commun. ACM* 51, 4 (April 2008), 34-40.
15. Sheng, S., Holbrook, M., Kumaraguru, P., Cranor, L. F., & Downs, J. (2010, April). Who falls for phish?: a demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 373-382). ACM.
16. Vishwanath, A., Herath, T., Chen, R., Wang, J., & Rao, H. R. (2011). Why do people get phished? Testing individual differences in phishing vulnerability within an integrated, information processing model. *Decision Support Systems*, 51(3), 576-586.
17. Wright, R. T., & Marett, K. (2010). The influence of experiential and dispositional factors in phishing: an empirical investigation of the deceived. *Journal of Management Information Systems*, 27, 273-303.
18. Wright, R., Chakraborty, S., Basoglu, A., & Marett, K. (2010). Where did they go right? Understanding the deception in phishing communications. *Group Decision and Negotiation*, 19(4), 391-415.
19. Wu, M., Miller, R. C., & Garfinkel, S. L. (2006, April). Do security toolbars actually prevent phishing attacks?. In *Proceedings of the SIGCHI conference on Human Factors in computing systems* (pp. 601-610). ACM.