Peering into the Phish Bowl: An Analysis of Real-World Phishing Cues

Lorenzo C. Neil[0009-0001-3084-7451], *North Carolina State University*
Shanée Dawkins[0000-0002-8114-0608], Jody L. Jacobs[0000-0002-6433-884X], Julia L. Sharp[0000-0002-9428-6518]
National Institute of Standards and Technology

Abstract

Organizations use simulated phishing awareness training exercises to help users identify, detect, and defend against the ever-changing phishing threat landscape. Realistic phishing emails are used to test users' ability to spot a phish from observable cues. However, much of the prior research does not focus on metrics for classifying the saliency of these cues. In this research, we analyzed different types of cues present in real-world phishing emails. The most common cues and cue types are presented, along with the frequency of their use in real-world phishing emails.

1. Introduction

Organizations within the public and private sectors implement embedded phishing awareness programs to prepare their employees for real-world phishing attacks. As a part of these programs, organizations send a simulated phishing email to employees to train them to spot a phish during low-risk exercises. The NIST Phish Scale (NPS) was created for these organizations to provide appropriate user context to their phishing awareness training data, typically click rates [1,2,3]. The NPS is a method for determining how difficult or easy a phishing email is to detect by considering both the observable characteristics of the email itself (cues) and the user context of the email's recipient (premise alignment). The research described here focuses on the first component, the cues.

The NPS identifies 23 cues that are categorized into five cue types. While the NPS helps cybersecurity practitioners identify which types of phishing emails are harder to detect within their organization, there appears to be a lack of research that creates a metric that not only aims at classifying the level of difficulty or identification of detecting phishing cues, but also that determines which cues are most salient (i.e., which cues are more identifiable and which are harder to detect than other cues). Prior work related to phishing cues focuses mostly on user detection towards specific types of cues or phish campaigns [4,5,6,7,8,9] while not addressing the saliency for multiple types of cues. Without such a classification, organizations face difficulties in assessing which cues impact their employees the most.

To this end, a research study was designed to investigate whether certain cues are more identifiable than others, and if certain cues are more likely to be ignored or not easily detected. To conduct this study, researchers first must determine the prevalence and frequency of cues in real-world phishing emails. Results of this initial examination by analyzing a compilation of real-world phishing emails are presented in this poster.

The findings from this work will inform a larger study to examine people's ability to identify different types of cues. Ultimately, this research will help organizations better tailor their phishing awareness training and help reduce an organization's security risk while still meeting their organization's mission and risk tolerance.

2. Method

Several collections of publicly available phishing emails (phish bowls) were researched for this investigation. The phishing emails contained in the phish bowls were actual emails reported to an organization's information technology office as suspicious and verified to be phishing attempts. One of the purposes of phish bowls is to assist end-users and help community members correctly identify suspicious versus safe and secure emails. Three phish bowls that included screenshots of the reported phishing emails were selected for this study (see Appendix section 7.2 for phish bowl considerations). These three phish bowls were from Taft College [10], the University of California, Santa Cruz [11], and the University of Vermont [12].

In this analysis, the NPS was applied to emails from the phish bowls. Certain properties of an email are required for this analysis; however, the screenshots from the phish bowls varied in consistency of providing the necessary information (e.g., missing email headers with sender and subject lines). Therefore, we used inclusion criteria to select a subset of the phishing emails for this study.

USENIX Symposium on Usable Privacy and Security (SOUPS) 2023. August 6 -- 8, 2023, Anaheim, CA USA.

In our phishing email dataset, 45 out of 82 total possible emails met our inclusion criteria, which required the screenshots to have email headers, subjects, sender display names, sender display addresses, and bodies of text. There were an additional 14 emails with these features redacted from the image by the phish bowl's organization. These emails were analyzed given the assumption that the inclusion criteria features were present in the original email. Assumptions made in analyzing these additional 14 emails were:

- All email and URL hyperlinks were not spoofed and directed users to the actual domain name or email address depicted in the email.
- Words in a text box denoting "Click here for file" are click-able and allow the end-user to download files.
- Attachments were present for phishing emails that referenced an attachment in the message body.
- Emails are customized to the targeted recipient's name and reference valid members of the organization.

Our total dataset from the phish bowls included 59 phishing emails, consisting of the 45 phishing emails that met the inclusion criteria, and the 14 emails where assumptions were made (see Appendix).

Three researchers independently rated the 59 emails with the NPS. Each email was then designated for two of the three researchers to compare individual NPS scores and come to a consensus NPS score for that email. The preliminary results for the cue analysis of phishing emails while applying the NPS reflect these consensus ratings for the phishing email dataset.

3. Results

This section provides a summary of the types of cues and the most common cues found in the phishing emails within our dataset. The five types of cues and their 23 associated cues from the NPS are listed in Table 1 in the Appendix.

3.1 Prevalence of Cues

The average number of cues in an email was 12.78 cues (n=59). The top two cues among phishing emails in our dataset were *Mimics a work or business process* (91.53%) and *Poses as a friend, colleague, supervisor, authority figure* (86.44%), both of which are Common Tactic cue types. Six of the top ten rated cues were either Language and Content (*Generic Greeting* – 79.66%, *Lack of Signer Details* - 61.02%, *Sense of Urgency* – 47.46%) or Technical Indicator (*Domain Spoofing* – 72.88%, *Sender Display Name and Email Address* - 47.46%, *URL Hyperlinking* - 42.37%) cue types.

We also calculated the total number of instances a single cue appeared in the 59 emails we studied. This calculation includes all instances a cue appears and not just the presence/absence count of a cue in each of the emails presented above. In total, we found 754 instances of cues present in the dataset. The total instance count for spelling and grammar irregularities topped the list of cues at 204 (27.06%). Five cues accounted for between five and ten percent of the total cues in the dataset (*Generic Greeting*, *Domain Spoofing*, *Mimics a work or business process*, *Poses as a friend, colleague, supervisor, authority figure*, and *Requests for sensitive information*). Fifteen cues accounted for less than five percent (see Appendix). Two cues were not present in the data (*Humanitarian appeals* and *Too good to be true offers*).

3.2 Cue Types

We find that of the five types of cues, four were present in at least 90% of the emails we analyzed (n=59): Language and Content (100%), Common Tactic (98.31%), Technical Indicator (94.92%), and Errors (93.22%). Visual Presentation Indicator (55.93%) was the least present cue type, appearing in just over half of the phishing emails.

Continuing the analysis of the total number cues present in the dataset (n=754), we examined these data by cue type. The Language and Content cue type had the highest total instance of cues in the dataset (30.77%). This was followed by Errors (30.11%), Technical Indicator (16.98%), Common Tactic (14.59%), and Visual Presentation Indicator (7.56%) cue types.

4. Conclusion

This work presents the first step towards investigating whether certain cues are more identifiable than others, and if certain cues are more likely to be ignored or not easily detected in phish training exercises. It is important to understand which cues are harder to identify by employees so organizations can properly perform and assess phishing awareness programs. To accomplish this, there needs to be a metric that classifies the level of detection difficulty in identifying cues, as well as determining the saliency of cues. Our findings from this retrospective observational study provide a depiction of the prevalence and frequencies of cues within real-world phishing emails from three publicly available phish bowls. The collected emails may not be representative of all phishing emails, rather are inherently the ones that have been detected. Yet, the sample of 59 emails is sufficient to inform our larger study in investigating people's ability to identify different types of cues.

5. Disclaimer

Any mention of commercial products or companies is for information only and does not imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the products are necessarily the best available for the purpose.

6. References

1. Greene, K. K., Steves, M., Theofanos, M., & Kostick, J.: User context: an explanatory variable in phishing susceptibility. In Proc. 2018 Workshop Usable Security (USEC) at the Network and Distributed Systems Security (NDSS) Symposium. (2018).

2. Steves, M. P., Greene, K. K., & Theofanos, M. F.: A phish scale: rating human phishing message detection difficulty. Proceedings 2019 Workshop on Usable Security. Workshop on Usable Security, San Diego, CA. (2019). https://doi.org/10.14722/usec.2019.23028.

3. Steves, M., Greene, K., & Theofanos, M.: Categorizing human phishing difficulty: A Phish Scale. Journal of Cybersecurity, 6(1), tyaa009. (2020). https://doi.org/10.1093/cybsec/tyaa009.

4. Valecha, Rohit, Pranali Mandaokar, and H. Raghav Rao. "Phishing email detection using persuasion cues." IEEE transactions on Dependable and secure computing 19.2 (2021): 747-756.

5. Moreno-Fernández, María M., et al. "Fishing for phishers. Improving Internet users' sensitivity to visual deception cues to prevent electronic fraud." Computers in Human Behavior 69 (2017): 421-436.

6. Dhamija, Rachna, J. Doug Tygar, and Marti Hearst. "Why phishing works." Proceedings of the SIGCHI conference on Human Factors in computing systems. 2006.

7. ThreatList: Top 5 Most Dangerous Attachment Types. https://threatpost.com/threatlist-top-5-most-dangerous-attachment-types/144635/. (Accessed August 2022).

8. Williams, Emma J., Joanne Hinds, and Adam N. Joinson. "Exploring susceptibility to phishing in the workplace." International Journal of Human-Computer Studies 120 (2018): 1-13.

9. Tsow, Alex, and Markus Jakobsson. "Deceit and deception: A large user study of phishing." Indiana University. Retrieved September 9 (2007): 2007.

10. Taft College Phish Bowl. https://ct-prod-wp.taftcollege.edu/its/phish-bowl-2/. (Accessed March 2023).

11. The University of California, Santa Cruz Phish Bowl. https://its.ucsc.edu/security/phish-bowl.html. (Accessed March 2023).

12. The University of Vermont Phishbowl. https://www.uvm.edu/it/phishbowl. (Accessed March 2023).

7. Appendix

7.1 Cues and Cue Types

Table 1 lists the five cue types and 23 cues that are components of the NIST Phish Scale [1,2,3].

Table 1: NIST Phish Scale cues

Cue Type	Cue
Error	Spelling and grammar irregularities
	Inconsistency
Technical indicator	Attachment type
	Sender display name and email address
	URL hyperlinking
	Domain spoofing
Visual presentation indicator	No/minimal branding and logos
	Logo imitation or out-of-date branding/logos
	Unprofessional looking design or formatting
	Security indicators and icons
Language and content	Legal language/copyright info/disclaimers
	Distracting detail
	Requests for sensitive information
	Sense of urgency
	Threatening language
	Generic greeting
	Lack of signer details
Common tactic	Humanitarian appeals
	Too good to be true offers
	You're special
	Limited time offer
	Mimics a work or business process
	Poses as friend, colleague, supervisor, authority figure

7.2 Phish Bowl Emails

There were multiple considerations that were involved in selecting the phish bowls for our phishing email analysis. These considerations included:

- Publicly accessible screenshots of phishing emails and not just descriptions or plain text renderings of the email's content.
- Visible headers (e.g., To:, From:, Subject:).
- Actual domain where embedded URLs are redirected (i.e., tooltips).
- Clear and concise embedded images and attachments.

The three phish bowls used were Taft College (Taft), University of California, Santa Cruz (UCSC), and University of Vermont (UV). Table 2 depicts the breakdown of the 59 emails from these phish bowls, including which emails met our inclusion criteria and which emails required specific assumptions to be made.

Table 2: Number of phish bowl emails included in the analysis

	Taft	UCSC	UV	Total
Number of emails that met inclusion criteria	22	23	0	45
Number of emails that required assumptions to be made	0	4	10	14
Total	22	27	10	59

Figures **Error! Reference source not found.** through **Error! Reference source not found.** are samples of the phishing emails that were used in our dataset.

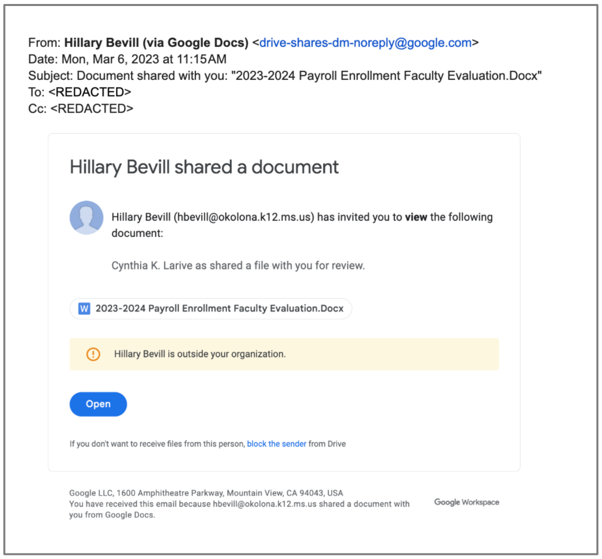

From: **Hillary Bevill (via Google Docs)** <drive-shares-dm-noreply@google.com>
Date: Mon, Mar 6, 2023 at 11:15AM
Subject: Document shared with you: "2023-2024 Payroll Enrollment Faculty Evaluation.Docx"
To: <REDACTED>
Cc: <REDACTED>

Hillary Bevill shared a document

Hillary Bevill (hbevill@okolona.k12.ms.us) has invited you to **view** the following document:

Cynthia K. Larive as shared a file with you for review.

W 2023-2024 Payroll Enrollment Faculty Evaluation.Docx

Hillary Bevill is outside your organization.

Open

If you don't want to receive files from this person, block the sender from Drive

Google LLC, 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA
You have received this email because hbevill@okolona.k12.ms.us shared a document with you from Google Docs.

Google Workspace

Figure 2. Sample file sharing phish type (UCSC phish bowl)

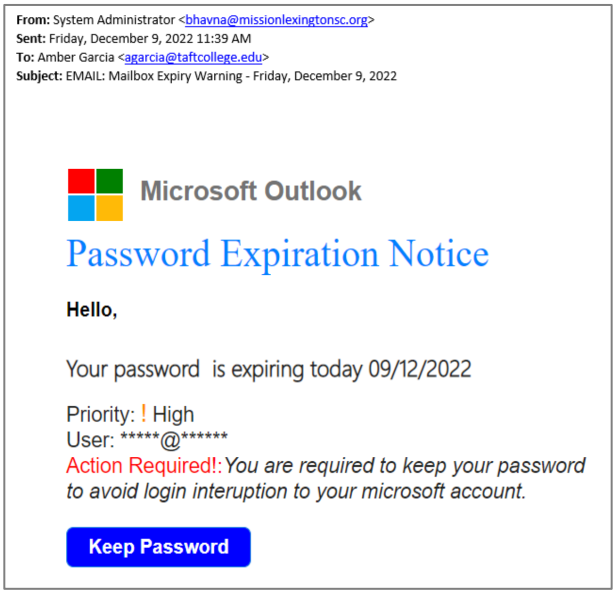

From: System Administrator <bhavna@missionlexingtonsc.org>
Sent: Friday, December 9, 2022 11:39 AM
To: Amber Garcia <agarcia@taftcollege.edu>
Subject: EMAIL: Mailbox Expiry Warning - Friday, December 9, 2022

Microsoft Outlook

Password Expiration Notice

Hello,

Your password is expiring today 09/12/2022

Priority: ! High
User: *****@******
Action Required!:*You are required to keep your password to avoid login interuption to your microsoft account.*

Keep Password

Figure 1. Sample password reset phish type (Taft phish bowl)

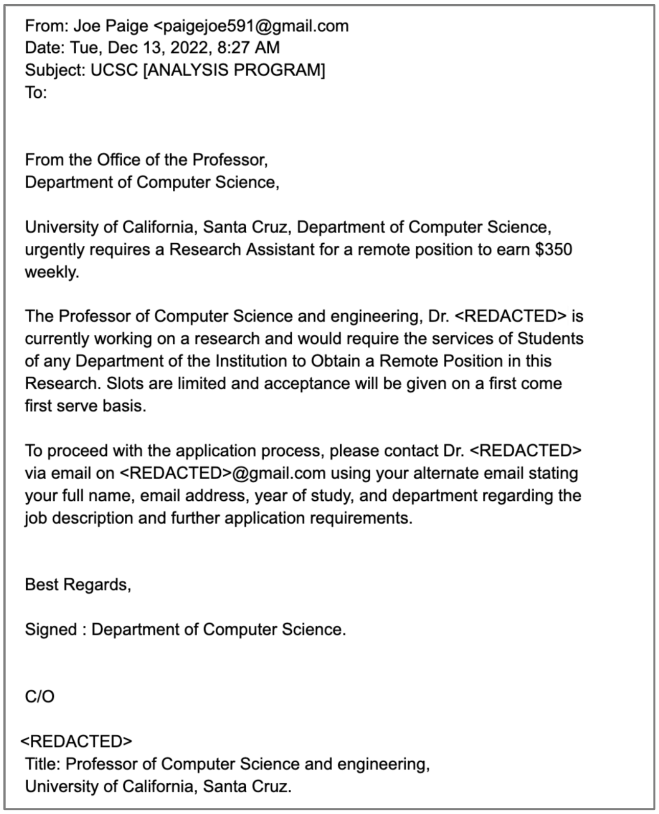

From: Joe Paige <paigejoe591@gmail.com
Date: Tue, Dec 13, 2022, 8:27 AM
Subject: UCSC [ANALYSIS PROGRAM]
To:

From the Office of the Professor,
Department of Computer Science,

University of California, Santa Cruz, Department of Computer Science, urgently requires a Research Assistant for a remote position to earn $350 weekly.

The Professor of Computer Science and engineering, Dr. <REDACTED> is currently working on a research and would require the services of Students of any Department of the Institution to Obtain a Remote Position in this Research. Slots are limited and acceptance will be given on a first come first serve basis.

To proceed with the application process, please contact Dr. <REDACTED> via email on <REDACTED>@gmail.com using your alternate email stating your full name, email address, year of study, and department regarding the job description and further application requirements.

Best Regards,

Signed : Department of Computer Science.

C/O

<REDACTED>
Title: Professor of Computer Science and engineering,
University of California, Santa Cruz.

Figure 3. Sample job announcement phish type (UCSC phish bowl)

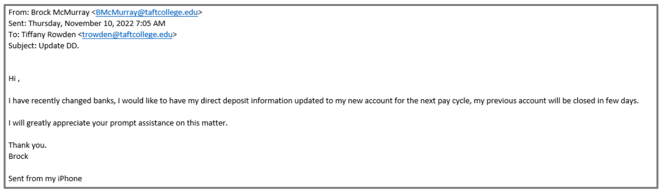

From: Brock McMurray <BMcMurray@taftcollege.edu>
Sent: Thursday, November 10, 2022 7:05 AM
To: Tiffany Rowden <trowden@taftcollege.edu>
Subject: Update DD.

Hi ,

I have recently changed banks, I would like to have my direct deposit information updated to my new account for the next pay cycle, my previous account will be closed in few days.

I will greatly appreciate your prompt assistance on this matter.

Thank you.
Brock

Sent from my iPhone

Figure 4. Sample direct deposit phish type (Taft phish bowl)

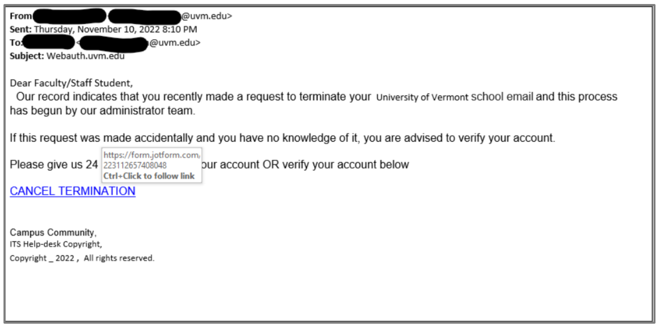

From [redacted] @uvm.edu>
Sent: Thursday, November 10, 2022 8:10 PM
To: [redacted] @uvm.edu>
Subject: Webauth.uvm.edu

Dear Faculty/Staff Student,
Our record indicates that you recently made a request to terminate your University of Vermont school email and this process has begun by our administrator team.

If this request was made accidentally and you have no knowledge of it, you are advised to verify your account.

Please give us 24 our account OR verify your account below

CANCEL TERMINATION

Campus Community,
ITS Help-desk Copyright,
Copyright _ 2022 , All rights reserved.

Figure 5. Sample account termination phish type (UV phish bowl)

7.3 Data Tables

This section presents our results from investigating the prevalence of cue types and cues within real-world phishing emails. Tables 3 and 4 list the counts and percentages of phishing emails with each cue and cue type present (n=59 emails). Tables 5 and 6 list the counts and percentages of the prevalence of cues and cue types in emails across the dataset (n=754 cues).

Table 3: Presence of cue types in phishing emails, ordered from cue type in the most emails to cue type in the least emails

Cue Type	Number of emails in which the cue type was present	% of emails in which the cue type was present
Language and Content	59	100.00%
Common tactic	58	98.31%
Technical indicator	56	94.92%
Errors	55	93.22%
Visual presentation indicator	33	55.93%

Table 4: Presence of cues in phishing emails, ordered from cues in the most emails to cues in the least emails

Cue	Cue Type	Number of emails in which the cue was present	% of emails in which the cue was present
Mimics a work or business process	Common tactic	54	91.53%
Poses as a friend, colleague, supervisor, authority figure	Common tactic	51	86.44%
Spelling and grammar irregularities	Errors	50	84.75%
Generic greeting	Language and content	47	79.66%
Domain spoofing	Technical indicator	43	72.88%

Cue	Cue Type	Number of emails in which the cue was present	% of emails in which the cue was present
Lack of signer details	Language and content	36	61.02%
Sender display name and email address	Technical indicator	28	47.46%
Sense of urgency	Language and content	28	47.46%
URL hyperlinking	Technical indicator	25	42.37%
Unprofessional looking or design	Visual presentation indicator	25	42.37%
Request for sensitive information	Language and content	20	33.90%
Inconsistency	Errors	17	28.81%
Legal language/copyright info/disclaimers	Language and content	13	22.03%
Threatening language	Language and content	13	22.03%
Attachment type	Technical indicator	9	15.25%
Distracting Detail	Language and content	8	13.56%
Logo imitation or out-of-date branding logos	Visual presentation indicator	7	11.86%
No/minimal branding and logos	Visual presentation indicator	5	8.47%
Security indications and icons	Visual presentation indicator	3	5.08%
Limited time offer	Common tactic	3	5.08%
You're special	Common tactic	2	3.39%
Humanitarian appeals	Common tactic	0	0.00%
Too good to be true offers	Common tactic	0	0.00%

Table 5: Prevalence of cue types in phishing emails, ordered from the most cue types across emails to least cue types across emails

Cue Type	Number of instances of each cue type across emails	% of instances of each cue type across emails
Language and Content	232	30.77%
Errors	227	30.11%
Technical indicator	128	16.98%
Common tactic	110	14.59%
Visual presentation indicator	57	7.56%

Table 6: Prevalence of cues in phishing emails, ordered from the most cues across emails to least cues across emails

Cue	Cue Type	Number of instances of each cue across emails	% of instances of each cue across emails
Spelling and grammar irregularities	Errors	204	27.06%
Generic greeting	Language and content	71	9.42%
Domain Spoofing	Technical indicator	57	7.56%
Mimics a work or business process	Common tactic	54	7.16%

Cue	Cue Type	Number of instances of each cue across emails	% of instances of each cue across emails
Poses as a friend, colleague, supervisor, authority figure	Common tactic	51	6.76%
Requests for sensitive information	Language and content	47	6.23%
Unprofessional looking design or formatting	Visual presentation indicator	36	4.77%
Lack of signer details	Language and content	36	4.77%
Sense of urgency	Language and content	35	4.64%
URL hyperlinking	Technical indicator	34	4.51%
Sender display name and email address	Technical indicator	28	3.71%
Inconsistency	Errors	23	3.05%
Threatening language	Language and content	18	2.39%
Legal language/copyright info/disclaimers	Language and content	16	2.12%
Logo imitation or out-of-date branding/logos	Visual presentation indicator	12	1.59%
Attachment type	Technical indicator	9	1.19%
Distracting detail	Language and content	9	1.19%
No/minimal branding and logos	Visual presentation indicator	5	0.66%
Security indications and icons	Visual presentation indicator	4	0.53%
Limited time offer	Common tactic	3	0.40%
You're special	Common tactic	2	0.27%
Humanitarian appeals	Common tactic	0	0.00%
Too good to be true offers	Common tactic	0	0.00%