User Experiences of TORPEDO: TOoltip-poweRed Phishing Email DetectiOn

Melanie Volkamer[a,c,1,*], Karen Renaud[b], Benjamin Reinheimer[c], Alexandra Kunz[c]

[a] *Karlstad University*
[b] *University of Glasgow*
[c] *Technische Universität Darmstadt*

Abstract

We propose a concept called *TORPEDO* to improve phish detection by providing just-in-time and just-in-place trustworthy tooltips. These help people to identify phish links embedded in emails. *TORPEDO*'s tooltips contain the actual URL with the domain **highlighted**. Link activation is delayed for a short period, giving the person time to inspect the URL before they click on a link. Furthermore, *TORPEDO* provides an information diagram to explain phish detection.We evaluated *TORPEDO*'s effectiveness, as compared to the worst case 'status bar' as provided by other Web email interfaces. People using *TORPEDO* performed significantly better in detecting phishes and identifying legitimate emails (85.17% versus 43.31% correct answers for phish). We then carried out a field study with a number of *TORPEDO* users to explore actual user experiences of *TORPEDO*. We conclude the paper by reporting on the outcome of this field study and suggest improvements based on the feedback from the field study participants.

Keywords: Phishing Detection, Email, Thunderbird, Usable Security, Tooltips, User studies

1. Introduction

Phishing is merely a modern equivalent of a confidence trick that has been carried out for centuries: deceiving someone to derive some personal benefit. The first time that the term "phishing" was used to refer to this digital version of confidence tricking was on January 2, 1996[1]. Phishing messages often offer a link embedded in an email message that entices the recipient to click. Email recipients are likely to click on links due to their widespread legitimate use. If they do click, it redirects them to a website masquerading as the real thing, or downloads malware onto their device.

*Corresponding author

Email addresses: melanie.volkamer@secuso.org (Melanie Volkamer), karen.renaud@glasgow.ac.uk (Karen Renaud), benjamin.reinheimer@secuso.org (Benjamin Reinheimer), alexandra.kunz@secuso.org (Alexandra Kunz)

[1] The mention occurred in alt.online-service.america-online.

Twenty years after its emergence, phishing still succeeds [12, 42]. Automated detection is a powerful tool against phishing, but the fact that it takes, on average, 28.75 hours to detect new phish websites [2] means that users have to detect phishing messages themselves during the discovery window. However, many people are unable to distinguish legitimate from phish messages. Since there is no financial bar on the number of emails phishers can send, this means a sizeable number of people are snared every day.

The goal of our research was to propose a way significantly to reduce phishers' success in the email environment (note, we studied the approach in the email environment but it could easily be adapted to other message formats such as Facebook and Twitter messages). To achieve this, we needed first to understand why people fall for phishing. We thus carried out a literature review and a cognitive walk-through analysis of emails as displayed by commonly-used desktop and webmail clients (Section 2). Based on our findings, we proposed a concept called *TORPEDO* (**TO**oltip-powe**RE**d **P**hish **E**mail **D**etecti**O**n) to assist users by providing a *just-in-time*, *just-in-place, trustworthy* tooltip that displays the actual URL with the domain *highlighted* in bold (see Fig. .1 for an overview and Section 3 for more details). Furthermore, we disable the link briefly, introducing a short delay, to increase the likelihood that people will check the URL before clicking on it. Finally, we provide users with an *extended information diagram* to explain the phish detection process (see Fig. .1). An evaluation of the proposal's effectiveness delivered significant improvements: 85.17% for phish, 91.57% for legitimate emails compared to 43.31% and from 63.66% when only providing the URL in the status bar (Section 4). The proposal was improved based on the results from this evaluation (Section 5). We implemented this improved proposal as a browser AddOn for Thunderbird, as a proof-of-concept. We subsequently carried out a follow-up study to assess user experiences and perceptions of the *TORPEDO* AddOn (Section 6). Most participants decided to retain the AddOn after the study, considering it useful in supporting phish detection. All participants provided helpful feedback to help us improve the AddOn even further. We discuss the feedback and our observations during our studies and propose the next iteration with an improved interface and extended functionality.

2. Why People Fall for Phish

It is important to understand why people fall for phish in order to be able to improve their resilience. To identify possible reasons for people falling victim to phish, we carried out a literature review and conducted a cognitive walk-through analysis.

2.1. Literature Review

A phishing email contains a number of signals that may indicate that the email is a phish, the most reliable of which is the actual URL as explained in [16] and [27]. Many people do not realise this but, just in case they do, phishers routinely obfuscate the URL to dampen down this signal (e.g., using `amazon.shop-secure.com` to phish amazon account holders). Our literature review revealed a number of papers in which the authors showed that many people focus on signals other than the URL, applying various flawed heuristics, namely:

- *The Sender:* People are likely to trust emails from friends [17] or from reputable businesses [44].

- *The Look and Feel:* People judge emails' trustworthiness based on their first impression, informed by a recognisable logo [5, 31], attractive design [31, 36], the use of their name or the provision of the company's contact details [18].

- *The Email Content:* People read the email in order to judge the trustworthiness thereof. Trusted indicators are the grammar and spelling quality [36, 31]. Researchers also showed that people are more likely to fall for a phish when: emotions such as excitement, fear or anxiety [5, 36] are elicited, a sense of urgency is invoked [31, 35], existing attitudes and beliefs (wanting to believe that the scammer is honest) are exploited [33], or persuasive and influencing techniques are used [36, 41]. Researchers argue that under such conditions of arousal people's decision-making abilities are impaired and they are less likely to detect and take note of danger signals [39].

- *Wrong Parts of the URL:* Some people do look at the URL [18]. However, they misinterpret the URL due to a lack of knowledge of the semantics of URLs [10, 43]. Some people are reassured by the mere presence of `HTTPS` in the embedded link and look no further [15]. Others are reassured due to the brand name being embedded '*somewhere in the URL*' [18].

2.2. Cognitive Walk-through Analysis

For one week, the authors carefully considered emails that we received, examining the embedded URLs to identify possible challenges which could impair or encourage phish detection. We examined Thunderbird and Apple Mail as well as Web interfaces from three popular Web email clients. We made the following observations:

- *Information not provided where expected/needed:* Thunderbird[2] as well as the Web email clients, display the actual URL destination in the status bar at the bottom of the window. *Problem:* The status bar is some distance away from the user's current attentional focus and might easily be missed. The text of the email is far more prominent and enticing and thus likely to be focused on.

- *Tooltip provided by sender:* Some email senders provide a tooltip which appears when the mouse hovers over the link. This happens in Thunderbird[3]. The actual URL is still displayed in the status bar. Providing such a tooltip is actually a very simple attack vector since the phisher only needs to override it by providing a "title=" attribute. *Problem:* The tooltip, which encourages examination of the URL by appearing where the user's attention is focused, actually potentially misleads because a phisher could easily provide a reassuring smoke-screen URL.

[2]Note, this is different for Apple Mail and also for Outlook.
[3]This is different for Apple Mail and for Outlook.

- *Redirection:* Some email providers seem to make phish detection difficult, if not impossible. These clients do not display the actual URL in the status bar, but rather an obfuscated URL, a so-called *dereferer* (see Fig. .2). Web mail providers claim that they do this to protect their users (due to carrying out checks before redirecting users to the actual web page).
  Problem: This approach makes it almost impossible for even the most security aware to detect the perfidy of the link.

- *Tiny URLs:* Some senders of legitimate emails use shortened URLs to redirect the person to a different website.
  Problem: Such shortened URLs make it impossible to detect the actual destination. It is advisable to use external services to find the final destination, but people do not necessarily know this.

- *Habituation:* While Apple Mail shows a toolbar next to the link in the email, it does so for both legitimate and phishing emails.
  Problem: Even if people know they should check the tooltip URL before clicking, they are unlikely to do this due to habituation or fatigue.

- *Mouse hover vs. clicking:* In order to get the relevant information in both desktop and Web email clients, one needs to hover over the link with the mouse — but not click.
  Problem: It is likely that users are not cautious enough and, instead of only hovering with the mouse, they click without checking (influenced by habitual actions).

2.3. Reasons Why People Fall for Phishing

Based on the observations reported in the previous section, the following reasons can be derived:

Reason 1: Not being aware that the URL is the only reliable signal: making a decision based on the wrong signal.

Reason 2: Not knowing which displayed URL to trust. There are three options: embedded in email text, in the displayed tooltip, or in the status bar.

Reason 3: Not having access to the actual URL (destination) due to URLs being obscured – either because of redirection or the use of tiny URLs.

Reason 4: Not consulting the URL carefully enough before clicking due to accidental clicks and/or habituation effects.

Reason 5: Not knowing how to distinguish authentic from phish URLs.

3. TORPEDO 1.0 as a Solution

We try to address all of the above-mentioned reasons with *TORPEDO*, which provides just-in-time and just-in-place trustworthy tooltips. These contain the actual URL with the domain **highlighted** (see Figure .3 for an example). It also delays link activation for a short period. Furthermore, *TORPEDO* provides an information diagram to explain phish detection, to be used together with the tooltips (at first use and subsequently on demand). We detail, in the following paragraphs, the different aspects and how they address the reasons identified in the previous section.

Just-in-time means that the tooltip appears when the person hovers their mouse over an embedded link. This addresses 'Reason 1' by making the reliable signal more prominent (at least compared to the status bar used by Thunderbird and Web email clients).

Just-in-place means that it appears right next to the link (i.e., right below the link). This addresses 'Reason 1' and 'Reason 2' by making it the most prominent signal and always displaying it in the same position, enhancing predictability.

Highlighting the domain in bold (similar to the highlighting in the address bar of some Web browsers) focuses attention on the most important part of the URL, addressing 'Reason 5'.

Disabling the link for a short period, perhaps three seconds, while providing continuous feedback in terms of a counter showing the time left to click (3, 2, 1s) increases the likelihood of people examining the link before clicking, addressing 'Reason 4'. Note, the delay is configurable to give users control. Furthermore, a white list is maintained to remember domains users have already clicked on twice before (requiring two clicks means that domains will not as easily be accidentally white listed). White-listed links will be activated immediately and not be subject to any delay so as not to annoy.

Trustworthiness, first, requires overwriting tooltips provided by the email sender. This, together with the tooltip appearing just-in-time and just-in-place, addresses 'Reason 2'. It also partially addresses 'Reason 3' by providing the actual URL, instead of the obfuscated one that the phisher would like the user to be fooled by.

Fig. .3 (a) shows how we propose to handle the redirections ('`redirectUrl=`') aspect of 'Reason 3', i.e., providing the actual URL and informing the user that this is the final destination. There are two possibilities for addressing the 'tiny URLs'[4] aspects of 'Reason 3': (1) automatically replace these URLs with the actual one using a service such as `http://www.wheredoesthislinkgo.com/`, or (2) inform users and let them decide whether to check for the actual URL using this service. From a usability point of view the first option looks more promising. However, from a security and privacy perspective the second one is safer (to implement the second the tool would have to send a requests even though the user has no intention of visiting the page). We decided to go for option (2) by default, but to allow users to configure for option (1). Thus, users would first see the upper tooltip in Fig. .3 (b) and the other one only once they have decided to check the actual URL.

[4] According to Wikipedia popular shortening services are: bit.ly, goo.gl, ow.ly, t.co, TinyURL, and Tr.im. The URL is parsed accordingly. Those services are addressed.

Information diagram, to support users in general, but in particular in checking the URL. This diagram (see Fig. .4) mainly addresses 'Reason 5'. It was refined and improved via several iterations based on feedback from laypeople. The diagram is shown when users initially start using the AddOn. Since users are not regularly confronted with phishing emails they may forget the rules after installation, so the diagram is also available on demand. The information diagram contains the following content, using a process approach to explain, step by step, what to do in terms of the URL manipulation tricks introduced by [40], namely obfuscating, misleading, mangling and camouflaging:

- In *Step 1*, we suggest focusing only on the URL. This addresses the fact that people do not focus on the URL ('Reason 1').The fact that the remaining parts of the URL can easily be faked is highlighted.
- In *Step 2* we recommend that people only focus on the highlighted part of the URL displayed in the tooltip.
- In *Step 3*, we explain that they should check whether the brand name is highlighted (to address misleading URLs such as `amazon.shop-secure.com` but also obvious phishes such as IP addresses). More precisely, we tell them to ignore the remaining parts of the URL.
- In *Step 4*, we advise them to check for extensions of the brand name such as in `amzon-shopping-in-America.com`. This the most difficult phishing URL to detect as some companies legitimately use such extensions in their authentic domains. We thus recommend that they search using Google if they are unsure.
- In *Step 5*, we recommend that they check, letter by letter, to identify small differences in the domain name (to pick up URLs such as `mircosoft.com`).

4. Evaluation of Effectiveness

We wanted to evaluate *TORPEDO*'s effectiveness, efficiency and user-engendered confidence in terms of properly judging the authenticity of emails, as compared to the *status quo* status bar display in Thunderbird and the considered Web email clients. To do so, we conducted a between-subjects online study launched on SoSciSurvey with participants randomly associated to one of two groups:

- **Status bar:** Sees the URLs in the status bar.
- **TORPEDO:** Sees the URL in a tooltip with domain highlighted in bold after having seen the information diagram.

Moreover, we formulated the following hypotheses:

- **H1 – Phish detection:** The *TORPEDO* group will detect more phishing emails, as compared to the status bar group.

- **H2 – Authentic eMail identification:** The *TORPEDO* group will identify more authentic emails, as compared to the status bar group.
- **H3 – Efficiency:** The *TORPEDO* group will judge emails more quickly, as compared to the status bar group[5].
- **H4 – Confidence:** The *TORPEDO* group will be more certain of their judgements, as compared to the status bar group.

4.1. Study Procedure

The study comprised the following three phases:

Phase 1 – Welcome: General information was provided, including the goal of the study, number of phases, the estimated duration, and data protection. We explained that it was important not to seek assistance (we did not elaborate by citing Google, as this could have been counter productive). We introduced the main tasks: They should imagine that their friend Max Müller is about to work through his emails. Since Max has accounts at all the companies in question, it is important for him to know which emails are authentic and which are phish. Therefore, they were asked to help him to judge the emails based on screenshots which Max provides to them on the following pages.

Phase 2 – Judging screenshots: Participants were presented with screenshots of 16 emails (8 authentic / 8 phish) each on a different web application and in random order. The *TORPEDO* group got in addition the information diagram, both before starting to answer questions as well as below each screenshot. Participants were asked: Is the email authentic? Then participants were then asked: How certain are you that you properly judged the displayed emails. The *TORPEDO* group was also asked to comment on the information diagram.

Phase 3 – Demographics: We requested demographic information.

Note, questions were translated from German to English for this paper.

4.2. Creation of Email Screenshots

We selected 16 web application providers based on the degree of popularity in Alexa (see Table .1). For all of these, we determined what authentic emails look like (including the 'from' address). All emails addressed the 'Heartbleed-Bug'. The text recommended that the recipient change their password and provided a link to facilitate this. The text differed slightly from one email to the next but the meaning remained the same. All exerted some (but not strong) pressure to change their password. Then, half of the screenshots

[5]Note, on the one hand it is important that people take their time to check the URL, however in addition, if they know what to consider, they can make decisions faster.

were 'turned into' a screenshot of a phish email by modifying the URL. For the *TORPEDO* group a tooltip was added to display the relevant link. We decided to simulate a worst-case scenario, i.e., advanced phishing emails which can only reliably be detected by checking the URL. We wanted to investigate the difference between the status bar and tooltip, and not the impact of various different signals. All emails were personalised. We also used `HTTPS` for both phish and non-phish displays because we did not want the absence or presence of `HTTPS` to constitute a cue (due to the findings in Section 2.1). Next, we considered which URL manipulation techniques to apply to obtain a representative set of manipulated URLs. Researchers have identified different URL manipulation classifications [40, 16, 26]. We used the categories from [40] with each type's anticipated success depending on how well users understand URLs and the thoroughness of their URL checking:

- *Obfuscate:* The phish URL is composed of an arbitrary name or IP address. Note, the brand name of the authentic website does not appear.
- *Mislead:* The phish URL embeds the authentic name somewhere (e.g., in the subdomain or the path) in order to allay suspicions.
- *Mangle:* The phish URL includes letter substitutions, different letter ordering, or misspelling e.g., `arnazon` instead of `amazon`.
- *Camouflage:* The domain name of the phish URL contains the brand name together with an extension or a different top level domain.

4.3. Ethics, Recruitment, and Incentives

Our University's ethical requirements with respect to respondent consent and data privacy were met. Participants first read an information page on which they were assured that their data would not be linked to their identity and that the responses would only be used for study purposes. Furthermore, using SoSciSurvey ensured that data was stored in Germany and thus subject to German data protection law. No debriefing was necessary. We recruited participants through a platform called Workhub, which is a German equivalent of Amazon Mechanical Turk. Every Workhub participant receives €3.

4.4. Results

The demographics are summarized for both groups in Table .2. Participants in the *TORPEDO* group, on average, detected phishing emails 85.17% of the time and they, on average, identified legitimate emails as such 91.57% of the time. The corresponding percentages for the control group are: 43.31% for phish detection and 63.66% for identifying legitimate emails. Note, participants, on average, detected Camouflaged URLs 72.09% of the time in the *TORPEDO* group and 38.37% in the status bar group, Misleading URLs 87.21% versus 30.23%, Mangled URLs 91.86% versus 37.20%) and Obfuscated URLs 89.53% versus 67.44%. Furthermore, the corresponding numbers for the answer 'I do not know' are (*TORPEDO*/status bar): 3.49%/6.98%, 2.91%/7.56%, 2.33%/8.14%, and 5.81%/5.23%. The descriptive data for H3 and H4 is depicted in Fig. .5.

As to the violation of homogeneity of variances for the compared groups we started our analyses with Mann–Whitney U tests for every hypothesis supplemented with an approximated effect size.

H1 – Phish Detection: Our analysis shows a significantly improved detection rate for phish Emails for participants in the *TORPEDO* group, as compared to those in the status bar group ($U = 210, p < .001$, $\eta^2 = 0.455$).

H2 – Authentic Email Identification: Our analysis shows a significantly improved identification rate of authentic Emails for participants in the *TORPEDO* group, as compared to those in the status bar group ($U = 304, p < .001$, $\eta^2 = 0.374$).

H3 – Efficiency: Our analysis shows that participants in the *TORPEDO* group were significantly more efficient than participants in the status bar group ($U = 676.5, p = .032$, $\eta^2 = 0.053$).

H4 – Confidence: Our analysis shows that participants in the *TORPEDO* group were significantly more certain about their decisions than participants in the status bar group ($U = 536.5, p < .001$, $\eta^2 = 0.155$).

4.4.1. Diagram Feedback.

We used open coding [37] to analyse the free text answers regarding the information diagram in Fig. .4. We came up with the following categories: 'grateful', 'nothing to improve', 'confusing', 'too much information', 'too little information', and 'small improvements'. Most of the participants (29 from 43 in total) were happy with the diagram, not mentioning anything to improve. Three were grateful. Eight mentioned that the diagram was confusing and five added that the confusion cleared once they read it. Two participants considered the diagram to contain too much information while another two participants complained about it giving too little information (requesting more examples). Three provided small suggestions for improvement: provide a title, and reconsider the usage of terms such as phish and URL.

4.5. Discussion.

The results show that, in the studied scenario, we significantly improved phish detection as well as the identification of legitimate emails with *TORPEDO*. The detection rates for all four phishing types increased, too. The participants in the *TORPEDO* group are also more confident that they made the proper decision in comparison to the status bar group. The decision making process is also significantly more efficient. Operating more quickly can, in the long run, lead to more errors being made. It is worth providing people with information such as that given in the information diagram since it is easy to apply and requires the email recipient to spend less time checking each individual email. The feedback to the information diagram showed that there is not much need to improve the diagram other than making the numbers clearer and adding a title. We acknowledge that the diagram might not provide

sufficient information for some people. In these few cases, we recommend extending our defence approach with existing proven training approaches (see Section 7).

Limitations. We acknowledge that we evaluated the approach in a best-case scenario as their primary task was security. Phishing effectiveness evaluations in field studies are not possible due to ethical and legal constraints. Lab studies also have their limitations because participants would not use a study laptop instead of their own. We also acknowledge that the URLs were displayed the entire time and not only when hovering the mouse over the link. This only partially simulates the proposed delay. Note, we only used `HTTPS` since we wanted to assess their ability to check the actual URL, not the presence or absence of `HTTPS`. Finally, we acknowledge that the sample was not representative.

5. TORPEDO 2.0

When we started implementing the AddOn, we considered how best to support future users. We incorporated lessons learned from the PassSec tool [40]. In effect, we decided to distinguish between the following three cases (mainly to avoid habituation). The tooltip frame is coloured in order to provide an additional signal about each individual case:

- *The low-risk case:* A green frame means the domain is classified as reliable based on a white list. The white list is based on the top 100 Alexa pages. In this case, the link is enabled immediately.
- *The uncertain case:* If the frame is red the domain should be checked carefully, because the link could be dangerous. In this case, the AddOn delays activation of the link for three seconds. Note that the settings allow users to customise the activation delay span. We decided to go for red and green despite some people being red/green color blind. This is not as big an issue as it could be since there are differences in the text as well in the activation delay (which only appears for the red case).
- *The track-record case:* The AddOn stores domains that the user has previously visited (derived by remembering the links they click on in emails). Once a domain has been visited at least twice, the domain has implicitly been deemed to be trustworthy by the user. In future, the tooltip will be orange for this domain and the link will be activated immediately. Note that the settings also allow the user to remove such domains from their personal white list.

Although the results of the first study suggested that mangled domains were not particularly challenging, we decided to present the domain name as proposed in [40], i.e.,

`a r n a z o n . c o m`

to maximize the support we provided. Although Thunderbird already warns users when the URL in the text does not match the actual URL in the link, we decided to treat this case, too, by adding a warning icon to the red case as well as the following statement: '*Warning! This could be a phishing attack, because the displayed domain differs from the actual domain.*'

Based on the results from the first study, we also refined the information diagram. The new information diagram is depicted in Fig. .1. Instructions to inform about the first step were added, wording was simplified and both formatting and style improved. Furthermore, the tooltips shown in the information diagram were also adapted. The tooltips for the shortened URL and the redirection URLs were adapted to be similar to the actual design of the main tooltip.

The AddOn is available from the Mozilla store[6] and has already been downloaded by over 1300 users. Furthermore, we developed an informative web page for *TORPEDO*[7] to introduce *TORPEDO*'s main functionality and to enlighten them about the nature of phishing. We will retain it to assist and inform future users. It is hosted at the University of Glasgow.

6. User Experience Field Study

We conducted a field study to evaluate user experiences of *TORPEDO*.

6.1. Methodology

The interaction with the participants proceeded either in person, via Skype or by telephone. The study was split into the following phases:

Phase 1 – Introduction: This phase welcomed participants and thanked them for their participation. We checked that they had installed and launched the **Thunderbird** email client and that they were connected to the Internet.

Phase 2 – Getting Started: We ensured that they understood the think-aloud process and communicated the fact that we would appreciate any feedback they provided. We informed participants that notes would be taken and requested their email address. Afterwards we collected basic demographics. This included how frequently they used Thunderbird daily, whether they could explain the term "Phishing", and asked how many such messages they received per week. We then briefly mentioned the name of the AddOn, and referred them to the `torpedo.website`. Participants were asked to read the information provided about *TORPEDO*.

Phase 3 – Installation: Participants were asked to install *TORPEDO*[8] and to re-launch Thunderbird. Once restarted, the *TORPEDO* welcome screen was displayed, followed by the information dialog shown in Fig .1. They were then asked three questions: (1) What do you like about the installation process? (2) What did you not like about it? (3) What would you recommend we improve?

[6] `https://addons.mozilla.org/en-GB/thunderbird/addon/torpedo-phishing-detection/`
[7] `http://torpedo.website`
[8] Note that `torpedo.website` also contains information about the installation process.

Phase 4 – Usage: Participants were referred to a newspaper story about people sending phish emails from other people's email addresses. We then explained that they were going to be asked to examine some emails from a friend called Max-Uwe. At least one of the emails was actually sent by a phisher. Their task was to judge the legitimacy of the emails.

Next we sent the participants four emails, from Max-Uwe's email address, in random order (see the Appendix for the actual text):

- One 'low-risk' email, i.e., resulting in a tooltip with green border;
- Three 'uncertain' emails, i.e., resulting in a tooltip with a red border. One is an easily detectable phishing URL, and another is a phishing URL that is hard to detect. The third is a trustworthy domain. Thus, two of the three should be judged as phish and one as authentic.

We asked them to work their way through the emails, commenting all the while and judging each. Afterwards, we asked them three questions: (1) What did you like about *TORPEDO*? (2) What did you dislike about *TORPEDO*? (3) What would you recommend we improve? We also asked them to fill in the SUS usability questionnaire [6].

Phase 5 – Post-Questions: We asked how likely they were to use the tool in the future, on a scale from 1 (extremely ulikely) to 5 (extremely likely). From those who wanted to retain *TORPEDO*, we requested permission to send them a link to a short survey a week later.

Phase 6 – Follow-up: A follow-up questionnaire was sent to participants who had expressed an interest in retaining the AddOn. The survey contained the following questions: (1) Did you change the configuration at all? If so, which settings and why did you change them? (2) What is the one thing you really liked about *TORPEDO*. (3) What is the one thing you really disliked? (4) What one thing would you advise us to change about *TORPEDO*?

6.2. Ethics, Recruitment, and Incentives

TU Darmstadt's ethical requirements were addressed as described in Section 4.3. Participants were recruited using social network posts, essentially implementing a snowball approach. Participants were told that they should use Thunderbird and that the goal of the study was to evaluate a Thunderbird AddOn which they will have to install and use during the study. Security was not introduced into the discussion, but it is likely that they made the connection. Participants did not receive any compensation. The interview took, on average, 35 minutes.

6.3. Findings of Pilot Study and Our Responses

The pilot study suggested two issues which we addressed by changing the design. First, participants were irritated by the 'i' icon in the tooltips, as shown in Figure .6 (a). The icon led them to believe that they would receive individualised information, which was not the case. We thus changed the icon to '?'. Furthermore, we explained, on the website, that they could access the information diagram via the '?'.

Second, participants were confused by the red and orange cases. They thought that the red signalled danger (despite explanations to the contrary on the website). Some even thought that orange was the one they ought to pay special attention to, where, in fact, we wanted red to signal that. Based on initial discussions with participants, we changed messages reflecting previous 'safe' decisions to blue. We also added instructions to the red tooltips telling them that they ought to check the highlighted part of the URL. Finally, we improved the explanation on the website.

6.4. Results

16 participants completed the first five phases and seven also completed the follow-up phase. Demographics for the entire group are depicted in Table .3.

6.4.1. Phase 1–5

The notes for the installation and usage phases were analysed independently by two of the authors. We took the notes from the think-aloud sessions as well as the responses to the specific questions.

For the *installation process* a number of positive aspects were identified from the notes. The process was considered to be easy, uncomplicated and fast. Explanations for this were that no additional data entry was necessary and no terms of use had to be read and accepted. Another positive aspect they referred to was that there was no need for additional packages or download managers to complete the installation. It was easily installed by using the standard Thunderbird functionality. Moreover, they mentioned that it was easy to decide which AddOn they had to install, because there was no other AddOn called *TORPEDO*. Some participants criticized that the information provided on the website about how to install the AddOn is not adequate for their browser-operation system combination. Note, some participants also commented on the website's content. The descriptions of phishing and the AddOn's functionality were described as easy to understand;

The *information diagram* was judged adequate and helpful by several participants. But some participants voiced some reservations:

- It was considered unnecessary, and they proposed only providing this information on request, e.g., by placing it in the settings area.
- It contained too much information on one screen (and it took them a while to parse it) and they proposed splitting it into several screens.
- It contained too much information to remember and recall. They proposed making this information available when the AddOn was used.
- It contained too little information, in particular they expected to see information from the *TORPEDO* website about the different cases. They proposed adding this extra information.
- It was not considered the appropriate media to convey the necessary information. They proposed using a different media type, perhaps a video-based tutorial.

- It was questioned whether the information diagram would be shown whenever Thunderbird launched (what one was afraid of).

For *using the AddOn to judge the four study emails* the following positive aspects were identified from the notes: *TORPEDO* is generally easy to use. They appreciated the use of different colors for different cases. Furthermore, the just-in-time and just-in-place information was mentioned. They were also positive about the fact that the link activation was delayed for three seconds. They felt this forced them to take the time to check the URL. Furthermore, they liked the highlighting of the domain and they also liked the the spacing between the letters in the domain. The following aspects were criticized by one or more participants. Some also proposed concrete improvements:

- There were no signs that *TORPEDO* was active after the installation. No icons were displayed to signal this. They had to hover over a link to confirm that it was executing.

- The text related to link activation was confusing in the German version of the tooltip. Note, that the English one is: '*Link is available in x seconds*'. while the corresponding German one translates to '*The link is activated in x seconds*'. Their interpretations were: (1) website will be opened automatically after three seconds; (2) more information will be provided after three seconds; (3) the time is an indication of how long the AddOn requires to check the trustworthiness of the URL.

- They expected to see context-specific information when they clicked on the '?' rather than seeing the same information diagram for all cases.

- It was too easy to lose focus while hovering over the link, with the consequence that the timer started again when the mouse was moved back over the link. A bigger "hitbox" for the tooltip was recommended or perhaps the AddOn could remember prior mouse hovering.

- While the information diagram recommended that they search for the domain if they were uncertain about the trustworthiness thereof, the process to do this is actually quite cumbersome. It was proposed that the tooltip itself provide direct access to this functionality, either by including an additional button or by allowing it via the right-mouse click options. One suggested that security websites such as 'virus total' be used to validate links rather than a search engine.

- While the track-record case was appreciated, they requested an option to directly inform the AddOn that a domain should be considered trustworthy.

One of the participants experienced difficulties distinguishing between green and red. He recommended the use of different border types (perhaps dashed) and that the instructions should refer not only to the red border but rather to the dashed red border.

One participant received the wrong links due to an error in the email set up. The email contained a redirection link. This meant we received some feedback about this tooltip too

(see Fig. .7). As this case was not explained on the website, the participant was confused by the button and it was necessary to explain it to him.

We also evaluated whether participants judged the four emails correctly. The low-risk email was correctly identified by all participants. All detected the typo in the domain of the phishing URL as well as the nonsense URL. The uncertain authentic email was problematic. A number participants who thought this one was phish used the fact that the tooltip border was red to confirm their initial hunch. One participant actually complained that *TORPEDO* was wrong in this case since it incorrectly judged the web address as phish. As a consequence, he wanted to uninstall the AddOn. It was necessary to provide further explanations.

Nine of the 16 participants said they were very *likely to use* TORPEDO *in the future*, and one participant was neutral. Three were unlikely to retain the tool. Those who went for either 'not' or 'not very likely' justified their decision by referring to their already heightened awareness and knowledge of phishing detection. They referred to the fact that the AddOn did not provide any useful additional functionality. However, they did consider the AddOn useful to others, like their parents. Twelve agreed to answer the follow-up survey after a week.

As a quantitative measure the System Usability Scale (SUS) describes a robust and versatile way to get user's subjective rating of a product's usability. Following over ten years of empirical evaluation there are various gradations to interpret the SUS scores e.g. a score of 85 and above could also be graded as an A [3, 6]. The SUS score was 86.25, which means that it achieved an excellent result in terms of usability.

6.4.2. Follow-up Phase

Two participants changed the settings: one disabled the link delay and the other one reduced the number of seconds. One participant wrote that he/she started recommending the tool to others. Furthermore, from the free text answers, we received the following additional feedback:

- Emails with many links (like ad-emails) are problematic. while moving the mouse over such an email, many tooltips appear and disappear again.
- Short URLs appeared and the appearance of the corresponding tooltip was confusing as it was neither mentioned in the study nor in the information dialog. There were questions related to the nature of the consequences.

6.5. Discussion and Derived Improvements

While we received positive feedback, we also gained new insights to help us to improve *TORPEDO*. The main problem on the conceptional level was that some participants interpreted the red border as a clear signal that the email was a phish. We wanted to communicate the fact that they ought to check the URL themselves since the AddOn was not able to verify the veracity of the URL. This is a crucial point as all authentic URLs that are not in the white list (and there will be many when they start using *TORPEDO*) are interpreted by these users as false positives. The users are likely to get irritated and uninstall the AddOn if they think it is being overly cautious (in their opinion) even though the URL is obviously

authentic. This needs to be avoided. Unfortunately, the information provided when clicking the '?' did not help to resolve their uncertainty as it only explained how to check the URL. Our proposal to address this problem includes making the following changes:

- The *low-risk case* gets an explanation, i.e., 'The developers classified the domain (highlighted part) of this web address as low-risk based on online white lists.' They can click on '?', to get more information about why the border is green and how the developers verified the authenticity of the domain.

- The *uncertain case* gets an explanation and the instructions are improved to 'Neither the developers nor you judged the domain (highlighted part) of this web address to be trustworthy. Please check the highlighted part of the web address carefully'. When clicking on '?', the user gets information on how to check a web address. Furthermore, the red border is replaced by a neutral black one.

- The *track-record case* gets an explanation, i.e., 'You previously classified the domain (highlighted part) of this web address as trustworthy'. When clicking on '?', the user gets more information about why the border is blue, i.e., that once a domain has been visited twice the AddOn adds it to the track-record white list. Furthermore, the user is told how to modify this list.

- The explanations and instructions for the special cases (short URL, redirections, and mismatch URL in text and actual URL) are also improved and the information when clicked on '?' adapted.

We plan to enable a deactivation of the explanation text in the tooltip if the user prefers smaller tooltips (by right clicking on the tooltip and selecting the corresponding item). With the added and adapted information being displayed when users click on '?', we also address the request for context-specific information.

Based on the input on the information screen, we decided to replace it with a tutorial with the following steps: (1) Welcome; (2) Provide general information on dangerous links and TOrPeDO; (3) Explain main cases including delayed activation; (4) Detail how to check the URL in the red case using three sub-screens (a) focus only on the highlighted part, (b) ensure that *only* the brand name is highlighted, (c) check for invalid extensions, (d) check for correct spelling; (5) Explain the special cases: red case, shortened URL and redirection; (6) Explain the features added to the tooltip menu; (7) Explain the additional adjustments that can be made via the settings. While we see the benefits of a professionally-produced video tutorial, the expense makes this untenable at present.

Minor improvements that we plan to implement are:

- Add the loop icon to Thunderbird's main frame to make it clear that *TORPEDO* is active.

- The text for 'Link available in 03 second(s)' will be changed to 'Link activation delayed to give you time to check it. Time remaining:'.

- Make the tutorial (the previous information diagram) optional, i.e., experts can skip this if they wish to.
- Explain, in the tutorial, how to find this information if there is a need to refresh their memory.
- Explain that the tutorial will not be displated at each launch of Thunderbird.
- Slightly extended hitbox. Note, here the challenge is to find the proper size. If the hitbox is too large two links next to each other might get mixed up.
- Slightly delay the tooltip appearance to avoid showing tooltips although the mouse moves over the link but does not stop there.
- When a user returns to a link after having seen a tooltip for that link, theactivation delay discounts the previous display time.
- Enable searching on a search engine for the domain when they right click on the tooltip and select the corresponding item.
- Allow white listing of a domain by means of a right click on the tooltip.
- Provide a new function in the AddOn settings to permit uploading of trustworthy domains.

Some suggestions we did not act upon. We decided to restart the timer once the tooltip disappears due to a focus shift for security reasons: the AddOn cannot distinguish between the user accidentally losing focus and focus being lost deliberately. Counting down while off-focus seems to defeat the purpose of having a timer at all. Furthermore, it was recommended to replace redirection URLs and short URLs immediately by the actual one. For privacy reasons, we decided not to do so by default. However, this can be activated in the settings.

Limitations: This study is more ecologically valid than survey-based studies because participants used their own laptops and their own email clients and accounts. However, when receiving the first four emails to check for phishing, they were informed by the study instructor that at least one of the four emails was a phish. The focus of this evaluation was the user experience as the effectiveness had already been evaluated in the first study. This limitation is thus acceptable. What needs to be acknowledged is that participants first read the website we provided about *TORPEDO* while people who download the AddOn from the Mozilla store have not necessarily read this additional information.

We acknowledge that most of the participants have a background in IT security and their performance, in terms of spotting phish messages, is probably a best case. Even so, we received valuable feedback and were able to improve and refine the AddOn.

7. Related Work

Researchers have proposed different ways of addressing phishing:

Automated Detection. Phishing emails can be detected either pre- or post-click. Emails can be analysed by the email provider before being forwarded to the user. This analysis includes checking the integrated URLs against several blacklists provided by companies such as Microsoft, Google and phishTank. Other checks can also be carried out. For example, to look at differences between displayed and actual domain names [13] or carry out an NLP analysis of the actual email text [38]. If the email is delivered and the person clicks, post-click detection can also occur. Web browsers or AddOns can check the URL against various blacklists or check the web site content in combination with the actual URL. A number of different approaches for these checks have been proposed [28, 4, 29, 32]. In both pre- and post-click checks a risky situation can either lead to blocking or a warning e.g., [25, 30, 45, 43]. As a final comment, there is an inherent flaw to post-click warnings. The human tendency to consistency makes it less likely for people to even want to detect the deceptive nature of any site if they have already committed to the process [9]. They have judged the email to be legitimate. Withdrawing at this stage is unlikely. *TORPEDO* does not aim to replace detection approaches but to complement them in order to help people to protect themselves in case none of the checks detects the phish or it is simply during the pre-detection window [2].

Training. A number of researchers have focused on training users to spot phish [1, 7, 8, 19, 21, 22, 23, 34] but most of them address phish detection in a web browser context. Researchers have shown that training improves phish detection rates. Training has two drawbacks as compared to *TORPEDO*. First, people need to be aware that there is a problem and that they need to be active in dealing with it. Otherwise they will not undergo training. This problem was addressed by Kumaraguru [24] by employing the concept of teachable moments, where people are given instructions or training when they almost fall for a phish. In their scheme, instead of blocking a link they allow it, and then show them that they almost fell for a phish. Second, people may forget the information the training imparted as they are not confronted with phishing emails every day. Again the teachable moment approach can help. However, we think providing the information graphic on demand, as and when required, is the safer approach as it might be that the next time they forget how to check the teachable moment mechanism may not be installed and they would be unprotected. Dodge *et al.* [11] report a different approach, post-click training. They send out fake phish messages and then train people who click on the links. They report a positive effect. However, this approach can only be taken within an organisation. A few participants in our evaluation had trouble understanding our diagram. For those, more exhaustive training may help them. Note, the training, as such, would be shorter than what would be required without *TORPEDO*.

Combining Approaches. We propose *TORPEDO* to complement existing approaches to address the email phishing problem. Other researchers have also proposed combining approaches to maximise phish protection. Khonji *et al.* [20] suggest a two-pronged approach, the first prong being user training, and the second being automated detection. The latter

includes blacklists, machine learning and visual similarity detection. Frauenstein and Von Solms [14] propose combining human, organisational and technical measures. The first includes awareness and training, the second policies and procedures and the last one includes automated measures to detect phish.

8. Conclusion and Future Work

Phishing is a thorny issue. Trying to filter out phishing emails before they reach the end users reduces the problem. Great strides have been made in this direction but no one will claim that any automated system will catch 100% of phishing emails. So, it is up to the email recipients to protect themselves. The research we have presented here offers a way to support end users by deploying *TORPEDO* to provide just-in-time, just-in-place, trustworthy tooltips. It also disables links for a short period of time, detects re-directions and tiny URLs, and provides a diagram that encapsulates phish detection advice in a step-by-step fashion. This approach has been iteratively evaluated and improved.

We found that it highly significantly improved phish (85.17% phish detection compared to 43.31% with the status bar URL) and legitimate email detection, made such detection significantly faster and led to people feeling more confident about their judgements as compared to the status bar approach. Based on the free text answers and the results for the different phishing types, we improved *TORPEDO* and developed a corresponding Thunderbird AddOn which is available in the Mozilla store. In a followup field study we studied the user experience and elicited input to further improve the AddOn. In particular, the appearance of the various tooltips will be improved.

As future work, we plan to implement *TORPEDO* for other web browsers by developing a Chrome AddOn. We will also consider mobile clients. Here, we discuss focusing on the grey case, i.e. not adding any action for blue and green but to do so for the grey case. In this case, a new dialog would be shown once the user clicks on the link. The dialogue recommends checking the displayed URL and then deciding whether to open the link or not. Similar to *TORPEDO*, by default the activation is disabled for 3 seconds.

In summary it can only be hoped that the different email clients and email providers deploy this approach as soon as possible as the disadvantage of an AddOn is that users need to be aware of its existence and actually install it.

Acknowledgement. This work was developed within the project 'KMU AWARE' which is funded by the German Federal Ministry for Economic Affairs and Energy under grant no. BMWi-VIA5-090168623-01-1/2015. The authors assume responsibility for the content. We would like to thank Peter Mayer for his support during the user study. Our special thanks go to Betty Ballin and Kristoffer Braun for the development of the AddOn and help improving the idea.

References

[1] A. Alnajim and M. Munro. An Anti-Phishing Approach that Uses Training Intervention for Phishing Websites Detection. In *6th International Conference on Information Technology: New Generations*, pages 405–410. IEEE, 2009.

[2] APWG Internet Policy Committee. Global Phishing Survey: Trends and Domain Name Use in 2H2013, 2013. http://docs.apwg.org/reports/APWG_GlobalPhishingSurvey_2H2013.pdf Accessed 13 March 2016.
[3] A. Bangor, P. T. Kortum, and J. T. Miller. An empirical evaluation of the system usability scale. *Intl. Journal of Human–Computer Interaction*, 24(6):574–594, 2008.
[4] Z. Bar-Yossef, I. Keidar, and U. Schonfeld. Do not crawl in the DUST: Different URLs with similar text. *TWEB*, 3(1):1–31, ACM, 2009.
[5] M. Blythe, H. Petrie, and J. A. Clark. F for fake: four studies on how we fall for phish. In *CHI*, pages 3469–3478. ACM, 2011.
[6] J. Brooke. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996.
[7] G. Canova, M. Volkamer, C. Bergmann, and R. Borza. NoPhish: An Anti-Phishing Education App. In *Security and Trust Management*, pages 188–192. LNCS, 2014.
[8] G. Canova, M. Volkamer, C. Bergmann, R. Borza, B. Reinheimer, S. Stockhardt, and R. Tenberg. Learn to Spot Phishing URLs with the Android NoPhish App. In *Information Security Education Across the Curriculum*, pages 87–100. Springer, 2015.
[9] R. B. Cialdini, J. T. Cacioppo, R. Bassett, and J. A. Miller. Low-ball procedure for producing compliance: commitment then cost. *Journal of Personality and Social Psychology*, 36(5):463, APA, 1978.
[10] R. Dhamija, J. D. Tygar, and M. Hearst. Why Phishing Works. In *CHI*, pages 581–590. ACM, 2006.
[11] R. C. Dodge, C. Carver, and A. J. Ferguson. Phishing for user security awareness. *Computers & Security*, 26(1):73–80, Elsevier, 2007.
[12] J.-P. Erkkilä. Why we fall for phishing. In *Conference on Human Factors in Computer Systems*. ACM, 2011.
[13] I. Fette, N. Sadeh, and A. Tomasic. Learning to detect phishing emails. In *16th International Conference on World Wide Web*, pages 649–656. ACM, 2007.
[14] E. D. Frauenstein and R. von Solms. Phishing: How an Organization can Protect Itself. In *Information Security South Africa Conference*, pages 253–268. Information Security South Africa, 2009.
[15] B. Friedman, D. Hurley, D. C. Howe, E. Felten, and H. Nissenbaum. Users' conceptions of web security: a comparative study. In *CHI*, pages 746–747. ACM, 2002.
[16] S. Garera, N. Provos, M. Chew, and A. D. Rubin. A framework for detection and measurement of phishing attacks. In *Recurring Malcode*, pages 1–8. ACM, 2007.
[17] T. N. Jagatic, N. A. Johnson, M. Jakobsson, and F. Menczer. Social Phishing. *Communications of the ACM*, 50(10):94–100, ACM, 2007.
[18] M. Jakobsson, A. Tsow, A. Shah, E. Blevis, and Y.-K. Lim. What instills trust? A qualitative study of phishing. In *FC*, pages 356–361. LNCS, 2007.
[19] K. Jansson and R. von Solms. Simulating malicious emails to educate end users on-demand. In *3rd Symposium on Web Society*, pages 74–80. IEEE, 2011.
[20] M. Khonji, Y. Iraqi, and A. Jones. Phishing detection: a literature survey. *Communications Surveys & Tutorials, IEEE*, 15(4):2091–2121, IEEE, 2013.
[21] I. Kirlappos and M. A. Sasse. Security Education against Phishing: A Modest Proposal for a Major Rethink. *Security and Privacy*, 10(2):24–32, IEEE, 2012.
[22] P. Kumaraguru, Y. Rhee, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge. Protecting People from Phishing: The Design and Evaluation of an Embedded Training Email System. In *CHI*, pages 905–914. ACM, 2007.
[23] P. Kumaraguru, Y. Rhee, S. Sheng, S. Hasan, A. Acquisti, L.-F. Cranor, and J. Hong. Getting Users to Pay Attention to Anti-phishing Education: Evaluation of Retention and Transfer. In *Anti-phishing WG*, pages 70–81. ACM, 2007.
[24] P. Kumaraguru, S. Sheng, A. Acquisti, L. F. Cranor, and J. Hong. Teaching Johnny Not to Fall for Phish. *Transactions on Internet Technology*, 10(2):1–7, ACM, 2010.
[25] L. Li and M. Helenius. Usability evaluation of anti-phishing toolbars. *Journal in Computer Virology*, 3(2):163–184, Springer, 2007.
[26] E. Lin, S. Greenberg, E. Trotter, D. Ma, and J. Aycock. Does domain highlighting help people identify

phishing sites? In *CHI*, pages 2075–2084. ACM, 2011.
[27] J. Ma, L. K. Saul, S. Savage, and G. M. Voelker. Beyond Blacklists: Learning to Detect Malicious Web Sites from Suspicious URLs. In *15th SIGKDD*, pages 1245–1254. ACM, 2009.
[28] S. Marchal, J. François, T. Engel, et al. Proactive discovery of phishing related domain names. In *Attacks, Intrusions, and Defenses*, pages 190–209. LNCS, 2012.
[29] M.-E. Maurer and D. Herzner. Using Visual Website Similarity for Phishing Detection and Reporting. In *CHI*, pages 1625–1630. ACM, 2012.
[30] M.-E. Maurer, A. D. Luca, and S. Kempe. Using data type based security alert dialogs to raise online security awareness. In *SOUPS*, page 2. ACM, 2011.
[31] R. Naidoo. Analysing Urgency and Trust Cues Exploited in Phishing Scam Designs. In *10th International Conference on Cyber Warfare and Security*, page 216. Academic Conferences Limited, 2015.
[32] P. Prakash, M. Kumar, R. R. Kompella, and M. Gupta. PhishNet: Predictive Blacklisting to Detect Phishing Attacks. In *INFOCOM*, pages 1–5. IEEE, 2010.
[33] J. J. Rusch. The "social engineering" of Internet fraud. In *Internet Society Annual Conference*. Internet Society, 1999.
[34] S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. F. Cranor, J. Hong, and E. Nunge. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *SOUPS*, pages 88–99. ACM, 2007.
[35] F. Stajano and P. Wilson. Understanding scam victims: seven principles for systems security. *Communications of the ACM*, 54(3):70–75, ACM, 2011.
[36] University of Exeter School of Psychology. The psychology of scams: Provoking and committing errors of judgement, University of Exeter, 2012.
[37] C. Urquhart. *Grounded theory for qualitative research: A practical guide*. Sage, 2012.
[38] R. Verma, N. Shashidhar, and N. Hossain. Detecting phishing emails the natural language way. In *ESORICS*, pages 824–841. Springer, 2012.
[39] A. Vishwanath, T. Herath, R. Chen, J. Wang, and H. R. Rao. Why do people get phished?Testing individual differences in phishing vulnerability within an integrated, information processing model. *Decision Support Systems*, 51(3):576–586, Elsevier, 2011.
[40] M. Volkamer, K. Renaud, G. Canova, B. Reinheimer, and K. Braun. Design and field evaluation of passsec: Raising and sustaining web surfer risk awareness. In *Trust and Trustworthy Computing - 8th International Conference, TRUST 2015, Heraklion, Greece, August 24-26, 2015, Proceedings*, pages 104–122, 2015.
[41] J. Wang, R. Chen, T. Herath, and H. Rao. An empirical exploration of the design pattern of phishing attacks. *Inform. Assurance, Security & Privacy Services*, Emerald Publishers, 2009.
[42] Webroot. Webroot 2015 Threat Brief. `http://www.webroot.com/shared/pdf/Webroot_2015_Threat_Brief.pdf` Accessed 13 March 2016.
[43] M. Wu, R. C. Miller, and S. L. Garfinkel. Do security toolbars actually prevent phishing attacks? In *CHI*, pages 601–610. ACM, 2006.
[44] Z. Xu and W. Zhang. Victimized by Phishing: A Heuristic-Systematic Perspective. *Journal of Internet Banking and Commerce*, 17(3), ARRAY Development, 2012.
[45] Y. Zhang, S. Egelman, L. F. Cranor, and J. Hong. Phinding Phish: Evaluating Anti-Phishing Tools. In *NDSS*. School of Computer Science, Internet Society, 2007.

Emails sent to Participants (translated from German)

E-Mail (green)

From: `max-uwe.maier@gmx.de`
Subject: Present for Tom's birthday

Hi,
As discussed, I've searched for different fishing tackle on Amazon. I have found two possible sets:

Set 1 for 22,89 Euro
Set 2 for 49,99 Euro

What do you think? We should discuss how much money we want to spend on it.

Best regards
Max-Uwe

`https://www.amazon.de/256tlg-Angelset-Teleskoprute-Angelspule-Angelkoffer/`
`...`

E-Mail (red, authentic)

From: `max-uwe.maier@gmx.de`
Subject: TU Dresden in university ranking

Hi,
As have you already noticed that TU Dresden did well in this ranking?

TU Dresden climbs in global ranking.

Best regards
Max-Uwe

`https://tu-dresden.de/tu-dresden/newsportal/news/the_ranking`

E-Mail (red, phish, easy to detect)

From: `max-uwe.maier@gmx.de`
Subject: Long weekend

Hi,
I have thought about where we could go on holiday. What about Venice? There is a good offer from ab-in-den-urlaub.

Best regards
Max-Uwe

`https://waser.ysed.ru/aerqerydx2342`

E-Mail (red, phish, hard to detect)

From: `max-uwe.maier@gmx.de`
Subject: New Media Market around your corner

Hi,
Have you noticed: the city is planning a new Media Markt around your corner:

Media Markt Opening

Best regards
Max-Uwe

`http://www.mediarnarkt.de/de/marketselection.html?label=darmstadt&lat=&lng=`

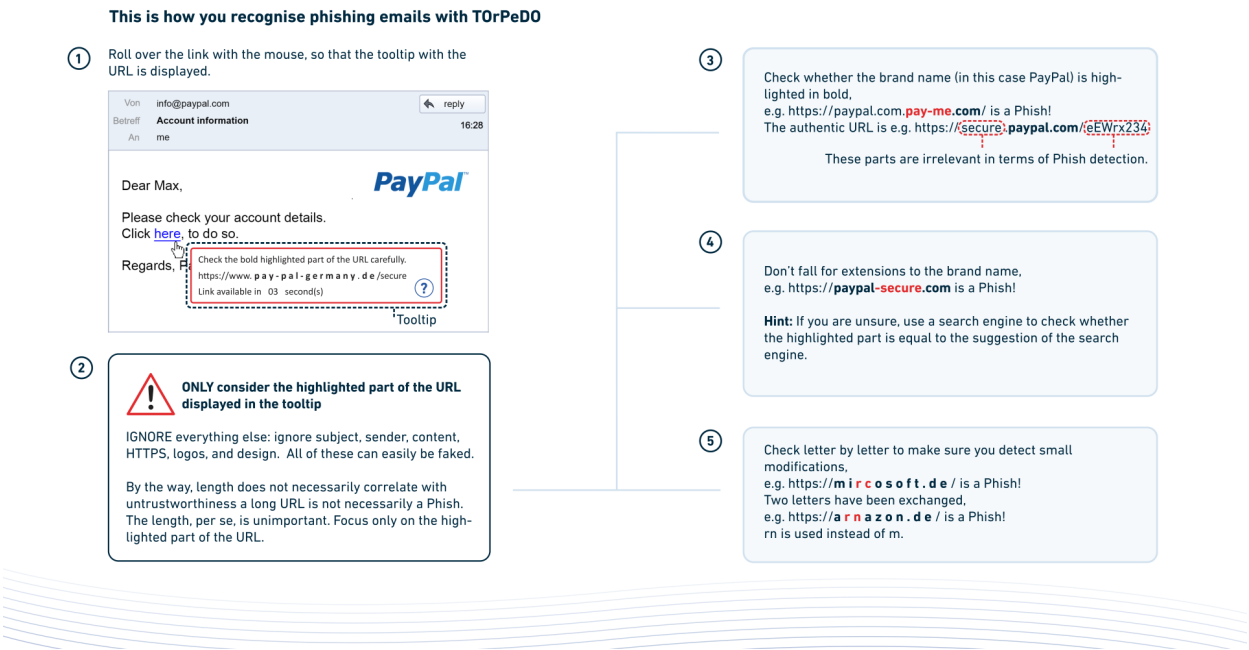

Figure .1: Just-in-time, just-in-place, trustworthy tooltips as shown in the upper-left part. The entire figure is displayed when starting the AddOn and when more information is requested.

Figure .2: Status bar displays an obfuscated URL for an embedded URL

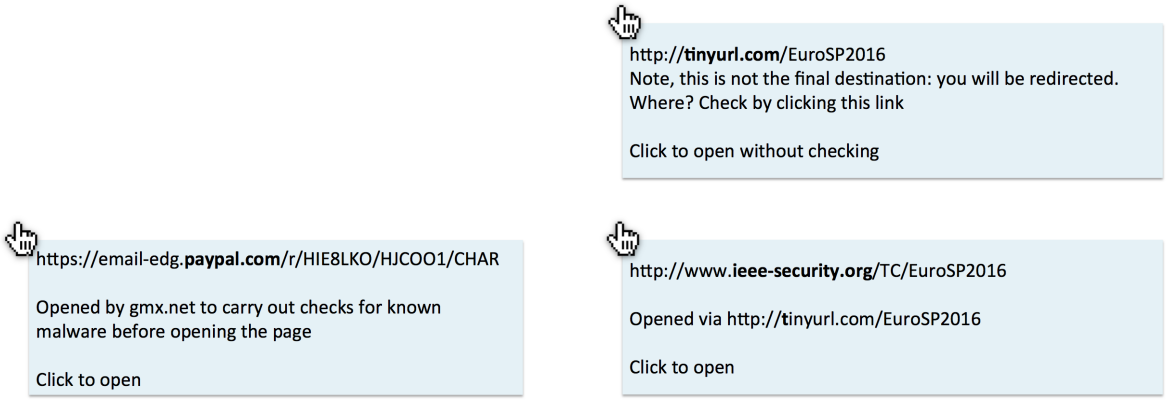

(a) Redirection case (b) Shortened URL case

Figure .3: Example tooltips - TORPEDO 1.0

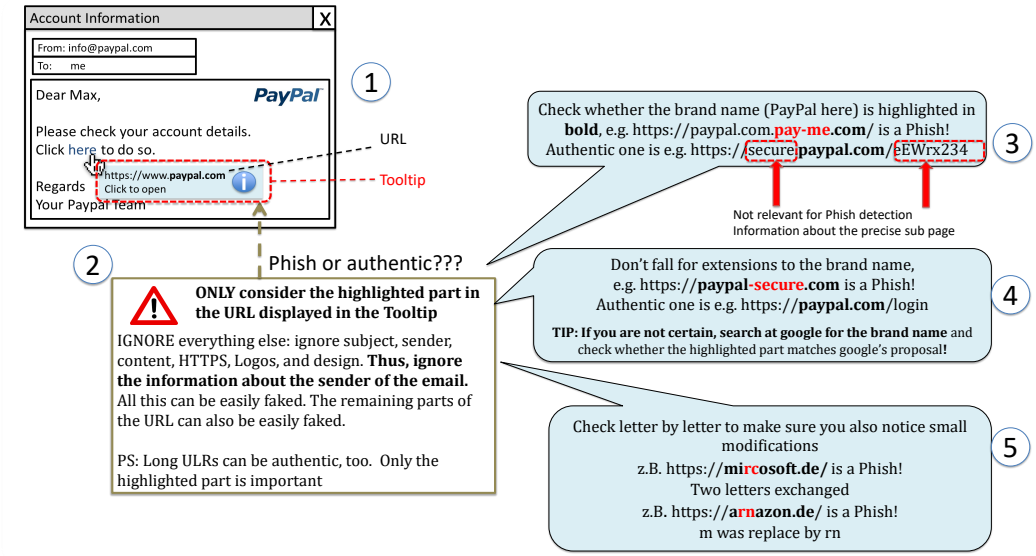

Figure .4: Information diagram - TORPEDO - 1.0.

Table .1: Legitimate(L) and Manipulated(M) URLs incl. type of manipulation.

Brand	**URL** (abbreviated with '...')
Postbank	L: `https://banking.postbank.de/rai/login`
Ebay	L:`https://signin.ebay.de/ws/eBayISAPI.dll?SignIn&UsingSSL...`
Xing	L: `https://login.xing.com/login?dest_url=https\%3A\%2F\%2Fwww...`
Google	L: `https://accounts.google.com/login?hl=de`
Dropbox	L: `https://www.dropbox.com/s/VPrize8EppElIxxWOwETRB87 Pe733AR...`
Telekom	L: `https://accounts.login.idm.telekom.com/oauth2/auth?response...`
Zalando	L: `https://www.zalando.de/login/`
MediaMarkt	L: `https://www.mediamarkt.de/webapp/wcs/stores/servlet/Logo...`
Facebook (Obfuscate)	L: `https://www.facebook.com/login` M: `https://130.83.167.26/login`
Flickr (Obfuscate)	L: `https://login.yahoo.com` M: `https://www.xplan.com/signing/flickr/`
Twitter (Mislead)	L: `https://twitter.com/login` M: `https://twitter.webmessenger.com`
Amazon (Mislead)	L: `https://www.amazon.de/ap/signin` M: `https://www.amazon.de.buecherkaufen.de/ap/signing?...`
DeutscheBank (Mangle)	L: `https://www.deutsche-bank.de` M: `https://meine.cleutsche-bank.de/trxm/db/i nit.do?login...`
Maxdome (Mangle)	L `https://www.maxdome.de/` M: `https://www.maxdorne.de/?fwe=true&force-l ogin-layer=true`
Paypal (Camouflage)	L: `https://www.paypal.com/signin/?country.x=DE&locale...` M: `https://www.paypalsecure.de/webapps/mpp/ho me`
GMX (Camouflage)	L: `https://www.gmx.net` M: `https://meinaccount.gmxfreemail.de/`

Table .2: Demographics

	# Participants	Average **Age**	Median	Youngest	Oldest	Male	IT Expert
Status bar	43	25.70	23	17	54	25	5
TORPEDO	43	27.86	26	18	60	25	2

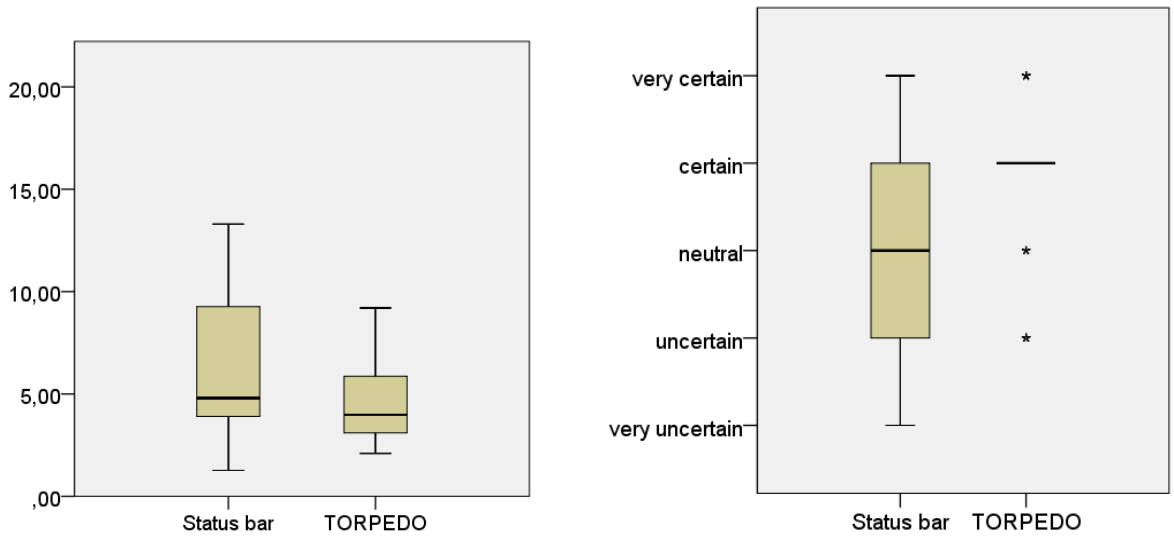

(a) Timing of correct decision (b) Certainty of correct decision

Figure .5: Descriptive data for timing and certainty of correct decision.

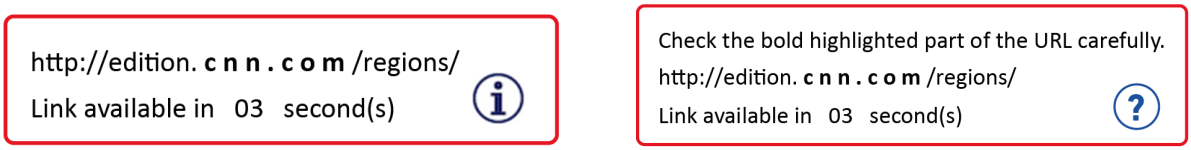

(a) TORPEDO 1.0 (b) TORPEDO 2.0

Figure .6: Tooltip for the uncertain case.

Table .3: Demographics

	# Participants
Gender	12 Male and 4 female
Age	18.75% under 30; 75% under 50; 33.33% over 50
IT Security Background	12 with security background; 4 without security background
Language Settings	11 German; 5 English
Reads Emails with Thunderbird	14 on daily bases - 2 weekly
Knows Phishing	15 yes; 1 no
How many Phishing Emails	3.02 per week

Figure .7: The tooltip for the redirect case.