Children's Understanding and Experiences with Online Phishing Threats

Tammes Burghard
Universität Paderborn

Abstract

As children are increasingly online, they become valuable targets for phishing scams. In order to protect their online safety, educating and enabling them to identify and avoid falling such scams is important. This research reviews existing literature on the matter and analyses children's understanding of phishing, their ability to identify and avoid it as well as educational approaches to enhance both of these skills.

It was found that even though children have some level of understanding of cyber security in general, they lack knowledge about phishing specifically. They also don't apply basic safety measures such as checking senders addresses or links. Various educational approaches to this have been explored with varying rates of success. Overall though, some promising improvements through education have been achieved being mostly held back currently by the lack of awareness, knowledge and confidence of teachers, schools and parents.

1 Introduction

According to Research by Kaspersky [3], over a third of children aged 11-15 years knows that they have been targeted recently by phishing attempts and more than a quarter has been a victim of a phishing scam. Children are a valuable target for online scammers because they make up a big portion of internet users, are relatively inexperienced and through them, their whole network of family and friends can be be attacked [5].

As automated means to combat phishing are insufficient, educating children about these threats, ideally before they fall to a fishing scam for the first time, is an important part of keeping them and their social surroundings safe online.

This paper will review existing literature in respect to the understanding and experience children have of phishing threats, their ability to identify and respond to phishing attempts and the role that education can play in protecting them.

Specifically, it will focus on the following research questions:

1. What theoretical understanding of phishing threats do children have?
2. How does this knowledge translate into their practical ability to identify and avoid phishing scams?
3. Which ways of educating children about phishing threats are there and how well do they work?

2 Methodology

In order to find papers to review, I used three methods:

1. Start with three papers provided by the seminar lead
2. Search ACM Digital Library and IEEE Xplore for the keywords "children" and "phishing"
3. Look up papers referenced in papers I had already found

The database searches provided 921 results for the ACM Digital Library and 47 results for IEEE Xplore. Based on their titles, most of the returned papers could quickly be disregarded as irrelevant for our research, because they either looked into phishing threats in general, adults understanding of them, or into general cyber security education. In the latter case, and if at least some focus on phishing seemed plausible, I checked the "Highlights" section of the ACM search result or searched the abstract (and in some cases the full text) for the term "phishing". I usually found that "phishing" was just used as an example cyber security thread that was not further looked into and therefore disregarded the paper. As they seemed to get less and less relevant, I also stopped checking ACM results after the first 400 entries.

Because this database search only gave me two relevant results in the end, I also tried to

- add "understanding" to the searched keywords
  - this provided basically the same results
- limit my search to title and/or abstract

    - this provided very few results that were a subset of the original results

- search for an entire research question

    - this provided very irrelevant results

In addition to the two relevant papers, I also found a Systematic Literature Review on children's general cyber security knowledge, skills and practice [4] where I was able to find relevant papers in its references.

Because I was only able to find five papers in total that focused on children's understanding of phishing threats, I decided to also include two more papers focusing on teenagers ability to detect phishing messages. Therefore, this literature review will look into the results of seven papers.

2.1 Reviewed Papers

The paper by Prior and Ophoff [7] covers a study that not only evaluates how well children recognize phishing in general, but also which specific concepts of phishing (outlined by the National Institute of Standards and Technology (NIST)) they know and are able to identify. The researchers analysed existing online resources and books for cyber security education with regards to which of the concepts of phishing they cover. Furthermore, the researchers did a training using one of those online resources in a UK primary school class with 43 children, 23 out of which gave consent to having their data analyzed in the study.

Steinmaurer et al. have developed the DigiSkill learning platform and in the paper [8] they evaluate its effectiveness in teaching Austrian students in the age of 12 to 17 about security in email communications and specifically identifying phishing emails.

Lastdrager et al. have done a comprehensive controlled study [5] in Dutch primary schools and analyzed 535 children's understanding of phishing threats with and without a short training session as well as how their understanding develops up to four weeks after the training. The study differentiates between the ability to detect legitimate emails in contrast to the ability to detect phishing scam.

The research done by Sun et al. [9] focuses on how children approach learning about phishing through an educational game and which types of activities they choose. With respect to these factors, the paper looks specifically into the impact of 'flow', which arises when the difficulty of a task is just right. For the study, 110 children in the age between 9 and 12 did a lesson in a computer lab. The game learning task too 30 minutes, another 45 minutes were spent on a pre-test, a post-test and a flow state questionnaire.

Alwanain did a study [1] researching children's ability to detect phishing messages on social media. He also designed an automated awareness training program using Whatsapp and examined its effectiveness in a controlled way by testing it on 30 children from Saudi primary schools.

For the study by Carle and Ophoff [2], a phishing education website for children was developed. It incorporated gamification and scenario-based learning by implementing a "choose-your-own-path style of game" [2]. It also avoided having overwhelmingly much text on the screen by splitting it into several slides that covered seven common phishing techniques. The study examined the effectiveness of visual cues that highlight important information for enhancing learning experience by having 18 children test the website either with or without such clues.

Nicholson et al. investigated teenagers' ability to identify phishing emails. In their study [6], 83 teenagers from North East England aged between 12 and 17 were tasked to evaluate 12 emails as either genuine or phish. The emails resembled the three common actions password reset, unauthorized access and email verification. They were all based on existing genuine emails by big internet services like Instagram or Netflix with the phishing ones having their sender address, name and links changed.

3 Results

While six of the seven studies measured children's phishing identification ability in order to evaluate an educational intervention, only two looked into their theoretical understanding of the matter upfront.

3.1 RQ1: What theoretical understanding of phishing threats do children have?

In the research by Prior and Ophoff, the 23 participating children self-reported to have "good knowledge of security [...] and how to identify a suspicious contact" [7]. However, they did not have a good understanding of phishing specifically and knew little about it compared to other cyber security topics like passwords and account protection. Only few of them knew what a phishing email is and no one had heard about smishing or vishing before. In addition to that, phishing did not seem like a primary focus for their parents either. They were more concerned about "1) access to inappropriate online content, 2) interacting with strangers, and 3) online bullying" [7] and did not consider email to be a particularly risky technology for their children.

Out of the 52 students that participated in the study by Steinmaurer et al. [8], 16 (30.07%) knew, what a phishing email is while all of them had an email address with the majority having multiple ones. In the pre-questionnaire including general cyber security questions, they showed a good performance with an average of more than 13 out of 15 correct answers. "32 participants (88.88%) answered that they would check a link by hovering over it before they click on it" [8] and almost all students knew that opening unwanted email attachments is unsafe. As typical characteristics of spam emails, they mostly highlighted "1) suspicious senders 2) suspicious links, and 3) spelling errors" [8].

3.2 RQ2: How does this knowledge translate into their practical ability to identify and avoid phishing scams?

Despite this decent theoretical understanding, the students in the study [8] struggled to apply it practically. Few of them actually checked a link by hovering over it before its opening and only 7 people (19.44%) thought an unwanted attachment to be suspicious. Instead, they "followed a strong visual approach" [8] and were much more likely to correctly identify phishing mails if they consist of plain text, maybe a logo and are not well laid out. If, however, the email is well designed, and does not seem dangerous on first glance, students blindly trusted it without applying the practices of caution that they theoretically knew of.

In the study done by Lastdrager et al. [5], the research team differentiated between children's ability to detect phishing and their ability to detect legitimate emails or websites. It found that they performed better at identifying phishing with an average score of 3.74 out of 5 compared to 2.26 out of 5 for legitimate questions. It further analyzed if there are predictor variables that influenced the children's score with the following results:

- sex had no significant effect
- older children scored higher than younger ones
- the school did have a significant effect
- children who had their own email address and/or their own Facebook profile scored higher than those without
- whether a child had already received a phishing email before or not did not influence its score significantly

Alwanain found in his study [1], that seven out of the 30 participating children opened the either malicious or legitimate link they were sent via Whatsapp. While gender did not have an effect on children's performance in identifying the phishing message, it did influence how they dealt with falling victim to it: "According to the parents' observations, all the male phishing victims blocked the sender immediately and reported the incident to WhatsApp, whereas the female victims informed their parents immediately after clicking on the suspicious link" [1].

Prior and Ophoff differentiated in their study [7] between eight specific concepts common to phishing emails that children where tasked to identify in 12 examples of phishing, smishing and vishing. The correct identification rate varied between the different concepts as shown in Figure 1.

Figure 1: Children's ability to identify specific concepts of phishing emails [7]

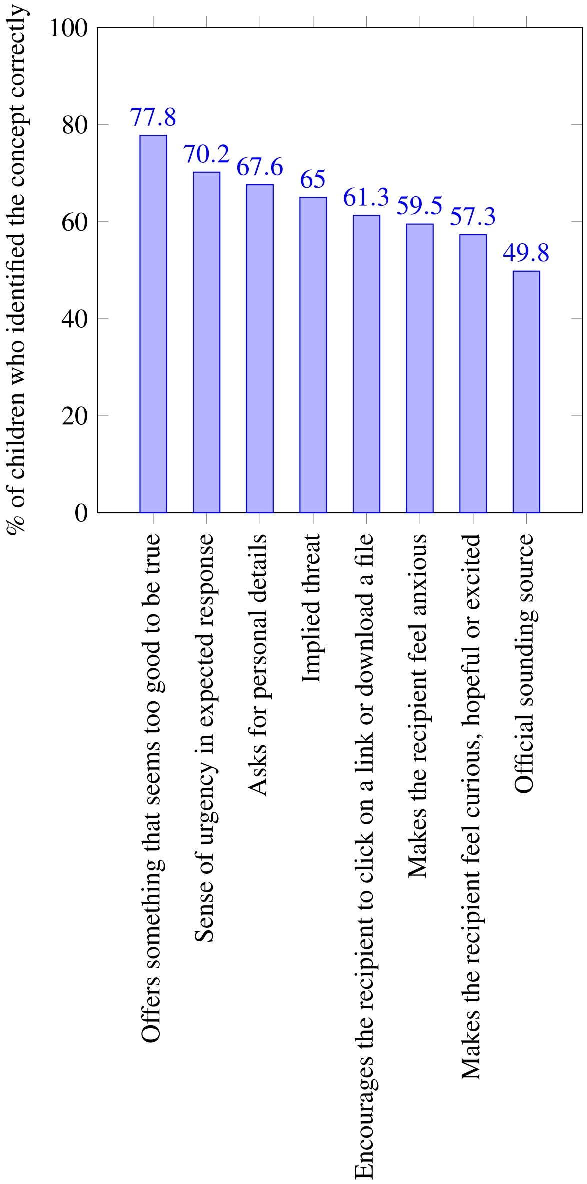

In the pretest of the study done by Carle and Ophoff [2], the 18 participating classified 49 out of 90 potential phishing emails correctly which lies with 54.44% just slightly above what would be expected from random guesswork.

In the study by Nicholson et al. [6], the participants identified 70% of the phishing emails correctly and 50% of the genuine emails with an overall success rate just below 60%. This performance was consistent throughout the age groups of 12-14 and 15-17 and it did not differ depending on whether the service the email pretended to come from was likely to be used by the participants or not. "[P]articipants were more accurate with messages focused on unauthorized account access as compared to password reset and account verification" [6] with a performance of 70% compared to between 53 and 55%. Given that all emails were designed to be easily classifiable by checking sender address and links, these techniques seem to not be used by teenagers even in a lab setting were they are "fully aware of the purpose of the experiment" [6].

3.3 RQ3: Which ways of educating children about phishing threats are there and how well do they work?

In the study by Steinmaurer et al. [8], the web based DigiSkill education platform was evaluated by doing interventions of 100 minutes length in computer science classes of Austrian secondary schools. Those interventions consisted of three phases:

1. Pre-phase
   - Pre-questionnaire
   - Task to classify five emails as spam or non-spam
   - IT security quiz covering 15 questions
2. DigiSkill intervention
   - prepared course in the DigiSkill tool
   - 30 minutes of time
   - following a fictive story
   - 13 tasks covering suspicious emails and websites as well as questions
3. Post-phase
   - Task to classify five different emails as spam or non-spam
   - IT security quiz covering 15 questions identical to the first quiz
   - Post-questionnaire

In the second quiz, students performance had improved considerably, but not significantly. This was probably due to a ceiling effect because the performance on the first test was already quite high not leaving much room for improvement.

For his study [1], Alwanain designed an automated education and test tool that sent 'malicious' links via WhatsApp to the participants of the treatment group and legitimate links to the ones of the control group. It then registered if the link was opened or not. The 'malicious' links lead to an educational website that informed the participants that they had fallen to a phishing attack and gave them information about common phishing techniques. In order to evaluate the learning effect, the experiment was repeated for a second time, now sending fake links to all participants. In total, the number of participants who opened the link dropped significantly from 7 out of 30 to 3 out of 30 between the two tests. However, no significant differences between the treatment group and the control group could be measured. According to Interviews that were done after the study, children's awareness for phishing threats had been positively influenced by "education that they had regularly received from parents and friends" [1] prior to the study.

Prior and Ophoff did a one hour lesson "based on CyberSprinters educational resources developed by the NCSC" [7], (the National Cyber Security Centre of the UK) in a UK primary school. The lesson started with a presentation introducing the children to the concepts of phishing, smishing and vishing. Here, different example emails were presented and the children voted "whether they thought it was genuine" [7] or not. Additionally, the reasons for voting one way or the other were discussed as well as the common features of phishing emails. After this presentation, 12 eample messages were placed around the classroom and the children were told that they now were part of a crime investigation tasked to identify all suspicious features of each message (for the specific features, see Figure 1). The children were excited for this task, but some lost focus at the end of the 20 minutes that were available for it. Most children "felt they had identified everything that indicated a scam, there was no consensus on any common theme as to how they achieved this." [7] In the end, the features were explained anew and the identification questions from the start of the lesson were repeated. The effectivity of this training was not evaluated.

The study by Sun et al. [9] did not focus on measuring and improving children's phishing identification performance. Instead, it uses an educational game in order to examine the relationship between flow experience, learning behavioral patterns and learning achievement. Flow experience was measured by asking the following two questions to learners: "(1) Do you feel that your skills were suitable for solving the tasks in this activity? and (2) How do you feel about the task's level of challenge in this activity?" [9]. Based on these answers, learners were grouped into the flow group if they rated their skill as appropriate for the games challenges, into the boredom group if the challenges were too easy and into the anxiety group if they were too hard. Based on these groups, the study analyzed, which of the four learning activities reading, interaction with peers, game and test the learners applied how much and in which order. It was found that the boredom group had much less peer interaction (2% instead of 13-14%) than the other groups and instead spent more time in the game dimension. They tended " to try the challenges repeatedly until they succeeded" [9] while the flow group combined independent learning with assistance from peers and then verifying "the information by reading the materials again" [9]. Learners in the anxiety group also learned independently and with peer assistance, but they did not verify what they had learned by reading again. However, the anxiety group was the only one that performed significantly better in the post-test compared to the pre-test.

The research by Carle and Ophoff "developed and evaluated a phishing education website for children aged 7–12 years old" [2]. The website uses several learning experience enhancing mechanisms such as visual cues, gamification, scenario-based learning and an avoidance of too much text being shown at once. The study specifically looked into the effectivity of visual cues that underlined important information in red in order to guide the attention to it by comparing between a test group that trained on a version of the website that included such cues and a control group using a website without them. In the post-test, the phishing identification performance had increased in both groups compared to the performance in the pre-test. Specifically, the test group went up from identifying 58% of the emails correctly to 73% while the control group increased their success rate from 51% to 67%. While these improvements are fairly similar between both groups, the test group enjoyed the experience more and found it easier to understand the content of the website, see Figure 2.

For their study, Lastdrager et al. [5] designed a training session that included a test and a 40 minute presentation. The presentation included topics like cyber bullying, hacking, phishing and identity theft and used storytelling with examples focussed on children. "Afterwards, [it] introduced four clues for identifying phishing emails: (1) how to find a URL from a hyperlink and how to assess where a URL leads to; (2) grammar, spelling, and the general type of language used; (3) presence of a sense of urgency or use of threats; and (4) the sender address. Furthermore, [it] showed two clues for websites: (1) the URL and (2) the need for an HTTPS connection when entering any data. During the training, the children were given ample opportunity to tell about their experiences" [5]. After the intervention, the children were tested by having them judge six emails and four websites. This was followed by a discussion of the correct answers. The control group just did the test without both the intervention and the discussion. After two, four, or 16 weeks, a similar test was repeated in order to measure knowledge retention. It was found that the training significantly improved the children's ability to detect phishing scam, but did not have a noticeable effect on their ability to detect legitimate emails. However, four weeks after the training intervention, this effect reversed. The children's ability to detect phishing went back to the level of the control group, but their ability to detect legitimate emails had improved significantly, see Figure 3.

Figure 2: Children's experience with the educational website [2]

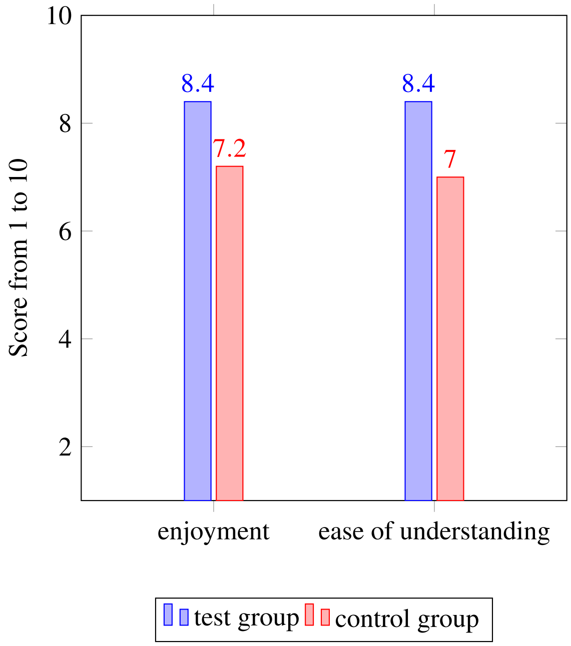

The effect of the simple awareness session in the study by Nicholson et al. [6] on the teenagers phishing identification performance has not been thoroughly evaluated. However, a feedback questionnaire with the participating teachers done three months after the initial session indicated that the session had long lasting positive effect on teenagers phishing identification confidence as well as their own confidence in teaching this topic. They also stated that this material should be included in the general curriculum and some of them had already suggested this content to others in their school.

4 Discussion

4.1 RQ1: What theoretical understanding of phishing threats do children have?

The fact that only two of the reviewed papers covered children's theoretical understanding of phishing, does not necessarily mean that this area has basically not being researched. There are plenty of publications about general cyber security education and it is likely that a subset of those does cover the theoretical understanding of phishing as a part of cyber security. What does become apparent though, is that phishing does not seem to be a primary focus not only for parents and

Figure 3: Knowledge retention [5]

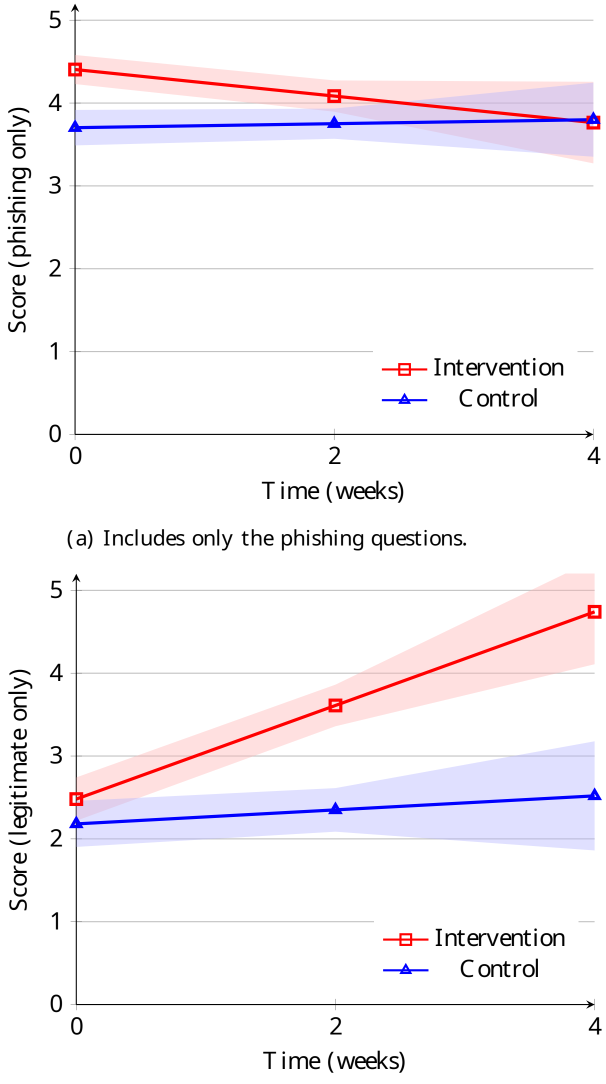

(a) Includes only the phishing questions.

(b) Includes only the legitimate questions.

education [7], but also in research. Both reviewed studies reported that only few young people knew what phishing is, let alone smishig or vishing [7, 8]. The results by Steinmaurer et al. hint that this lack of knowledge is only of terminological nature and the participants did have a decent understanding of the concepts of phishing. This finding remained fairly vague though, is not supported by other work that was reviewed here and needs further research.

As shown in RQ3, children's understanding of phishing can be effectively improved through education. As hinted in the paper by Prior and Ophoff, the main hurdle here seems to be that both parents and teachers don't have this topic on their radar and teachers lack the knowledge and confidence to teach this topic. Therefore, it is important to have "targeted education and training of teachers [...], especially for those outside computer science." [7]

4.2 RQ2: How does this knowledge translate into their practical ability to identify and avoid phishing scams?

Apart from Alwanain [1], all the reviewed research measured the ability to identify phishing emails or websites in a lab setting where the participants knew that their task was to identify phishing. This might explain, why Nicholson et al. [6] and Lastdrager et al. [5] found the performance of identifying phishing to be higher than for identifying legitimate content.

Despite these studies all measuring phishing identification performance in a similar way, they are not really comparable when it comes to which factors influence said performance. The participants of their studies are of different age and different countries and, in regard to this question, the studies focus on different things such as phishing techniques [7], preconditions of the participants [5], how the participants approach phishing identification [8], the type of phishing message and whether its sender is one that the participants regularly interacted with [6], or the effectiveness of an educational intervention [1, 2]. While this diversity in approaches nicely demonstrates the complexity of the matter, it makes it hard to draw general conclusions.

However, one important finding shared between both Steinmaurer et al. [8] and Nicholson et al. [6] is that participants did not check sender addresses or links before opening them. This highlights the need for education in this regard, but also for email clients that make it easy to check for these things for example by exposing the sender address prominently and adding a confirmation dialog to the link opening process.

4.3 RQ3: Which ways of educating children about phishing threats are there and how well do they work?

All the reviewed papers evaluated mostly self developed concepts for phishing education. Because these concepts are quite diverse and they are also tested on children of differing age, country and socialization, they provide a variety of options for phishing education that can only be compared to a very limited degree by comparing how much phishing identification performance of children has improved through the respective intervention. Additionally, as there is no standardized way to measure phishing identification performance, those results have to be taken with a big grain of salt as well.

Four studies evaluated digital educational tools, three of which are web based. Alwanain [1] developed a simple web page aimed at increasing awareness for the topic of phishing. The DigiSkill education platform evaluated by Steinmaurer et al. [8] is a flexible and modular learning management system and the website developed by Carle and Ophoff [2] featured several mechanisms that improved the learning experience such as visual cues, gamification, scenario-based learning and evaluated specifically the effectiveness of visual cues. Sun et al. [9] analyzed children's learning behaviour in a digital educational game.

In contrast to that, Prior and Ophoff [7] as well as Lastdrager et al. [5] evaluated classroom sessions featuring a presentation and interactive elements such as discussions and role play.

Of the digital tools, the one by Carle and Ophoff [2] was the only one that significantly improved children's phishing identification performance. I the game by Sun et al. [9] one of the three investigated groups significantly improved their scores. The tool by Alwanain [1] only aimed at improving awareness which seemed to have already been quite high and the intervention by Steinmaurer et al. did lead to considerable improvement, but the significance probably fell victim to a ceiling effect because the performance on the pre-test was already quite high not leaving much room for improvement in the post-test.

Of the classroom sessions, only Lastdrager et al. [5] properly evaluated its effectiveness including knowledge retention over a course of four weeks. According to them, the intervention initially significantly improved the children's ability to detect phishing. But after four weeks, this effect was gone and instead, the ability to detect legitimate content had improved significantly in the group that received the intervention, see Figure 3.

Prior and Ophoff [7] made the notable observation that when discussing phishing, children used examples that made more sense to them than what the material had provided. Therefore, they suggest that children's perspective should be taken into account when developing education material.

5 Conclusion

Overall, this research shows that phishing is a topic that often gets overlooked in cyber security education and not deemed as important by most parents. Even though children do have some level of understanding of cyber security in general, their knowledge regarding phishing specifically is fairly poor. And then even if they have some basic knowledge about it, they often fail to apply it when practically confronted with emails or other messages to judge. Especially, they don't check sender addresses and links before opening them.

Several approaches to improve this situation through educating children have been explored, all either in the form of digital educational tools or classroom sessions. The success of these approaches varied from study to study. But overall, positive impact could be achieved and participating teachers and students evaluated phishing education as helpful and important afterwards. The biggest hurdle seems to be the lack of knowledge and confidence of teachers highlighting the need for specifically educating them.

References

[1] Mohammed Alwanain. How do children interact with phishing attacks? 21:127–133, 03 2021.

[2] Nicole Carle and Jacques Ophoff. Using visual cues to enhance phishing education for children. In Lynette Drevin, Wai Sze Leung, and Suné von Solms, editors, *Information Security Education - Challenges in the Digital Age*, pages 21–35, Cham, 2024. Springer Nature Switzerland.

[3] Kaspersky Lab. Overconfident and over exposed: Are children safe online? 2023.

[4] Maria Lamond, Karen Renaud, Lara Wood, and Suzanne Prior. Sok: Young children's cybersecurity knowledge, skills & practice: A systematic literature review. In *Proceedings of the 2022 European Symposium on Usable Security*, EuroUSEC '22, page 14–27, New York, NY, USA, 2022. Association for Computing Machinery.

[5] Elmer Lastdrager, Inés Carvajal Gallardo, Pieter Hartel, and Marianne Junger. How effective is anti-phishing training for children? In *Proceedings of the thirteenth Symposium on Usable Privacy and Security, SOUPS 2017*, pages 229–239. USENIX Association, 2017.

[6] James Nicholson, Yousra Javed, Matt Dixon, Lynne Coventry, Opeyemi Dele Ajayi, and Philip Anderson. Investigating teenagers' ability to detect phishing messages. In *2020 IEEE European Symposium on Security and Privacy Workshops*, pages 140–149, 2020.

[7] Suzanne Prior and Jacques Ophoff. Lessons learnt from using educational phishing materials with uk primary

school children. In Lynette Drevin, Wai Sze Leung, and Suné von Solms, editors, *Information Security Education - Challenges in the Digital Age*, pages 36–49, Cham, 2024. Springer Nature Switzerland.

[8] Alexander Steinmaurer, Azra Bajramovic, Daniel Pollhammer, and Christian Gütl. Learning security awareness in email communication using a platform for digital skill teaching. In *2022 IEEE International Conference on Teaching, Assessment and Learning for Engineering (TALE)*, pages 38–45, 2022.

[9] Jerry Chih-Yuan Sun, Cian-Yu Kuo, Huei-Tse Hou, and Yu-Yan Lin. Exploring learners' sequential behavioral patterns, flow experience, and learning performance in an anti-phishing educational game. *Journal of Educational Technology and Society*, 20(1):45–60, 2017.