

# The Soft Skills of Software Learning Development: the Psychological Dimensions of Computing and Security Behaviours

Matthew R Ivory

Supervised by Prof. J Towse, Prof. M Levine, Dr. M Sturdee, & Prof. B Nuseibeh

Lancaster University

Lancaster, UK

matthew.ivory@lancaster.ac.uk

## ABSTRACT

When writing software code, developers typically prioritise functionality over security, either consciously or unconsciously through biases and heuristics. This is often attributed to tangible pressures such as client requirements, but little is understood about the psychological dimensions affecting security behaviours. There is an increasing demand for understanding how psychological skills affect secure software development and to understand how these skills themselves are developed during the learning process.

This doctoral research explores this research space, with aims to identify important workplace-based skills for software developers; to identify and empirically investigate the soft skills behind these workplace skills in order to understand how soft skills can influence security behaviours; and, to identify ways to introduce and teach soft skills to computer science students to prepare the future generation of software developers.

The motivations behind this research are presented alongside the work plan. Three distinct phases are introduced, along with planned analyses. Phase one is currently in the data collection stage, with the second phase in planning. Prior relevant work is highlighted, and the paper concludes with a presentation of preliminary results and the planned next steps.

## CCS CONCEPTS

• **Security and privacy** → **Human and societal aspects of security and privacy**; • **Software and its engineering**; • **Human-centered computing** → **Human computer interaction (HCI)**;

## KEYWORDS

Soft Skills, Cognitive Psychology, Security, Behavioural

### ACM Reference Format:

Matthew R Ivory. 2022. The Soft Skills of Software Learning Development: the Psychological Dimensions of Computing and Security Behaviours. In *The International Conference on Evaluation and Assessment in Software Engineering 2022 (EASE 2022), June 13–15, 2022, Gothenburg, Sweden*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3530019.3535344>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*EASE 2022, June 13–15, 2022, Gothenburg, Sweden*

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9613-4/22/06...\$15.00

<https://doi.org/10.1145/3530019.3535344>

## 1 INTRODUCTION

In this research, soft skills are defined as the psychological dimensions, or traits, that underpin behaviour [5]. Commonly, soft skills are synonymous with workplace-relevant transferable skills, including skills such as "teamwork" or "time management", and have been the focus of previous human factors research [10, 14, 15]. This current research seeks to go beyond these surface level traits, to identify the psychological dimensions that underpin transferable skills. In this body of research, transferable skills are referred to as "shallow skills", and soft skills are the underlying psychological dimensions. Shallow skills are referred to as such, because they provide little in the way of quantifiable skills, and their definitions often change depending on research context. Shallow skills can be considered as the manifestation of soft skills, particularly in workplace situations. Soft skills are the more immutable, psychological aspects of behaviour.

Software development is the direct product of human interaction, created through the combination of cognitive abilities, social interactions and the unique culture of software development [3, 23]. In recent years, the software industry has become aware of the significance of soft skills for successful software creation [5, 14, 15]. By 2030, there is an anticipated 22% increase in employment opportunities for software developers compared to an average 8% increase across all other industries<sup>1</sup>, but a rising concern that graduates entering the workforce are lacking the necessary cognitive and social skills required for successful integration into the workplace [13]. This issue has been evidenced in software security roles, with research indicating the most important skills required for security roles are not technical in nature, but are soft skills [8]. As a consequence, it is vital to identify the psychological traits required to successfully develop secure software.

Security in software is not a new concern, but the responsibility for security has changed over time. In 1999, Adams and Sasse [2] argued that software users were "not the enemy" and their fallible security behaviours were not their fault, but rather that of developers disregarding default user behaviour.

Similarly in 2008, Wurster and van Oorschot [24] posited that developers were "the enemy" and as they are the ones causing security issues, security should be removed from their responsibilities. They suggested the onus should be placed with API developers as they provide functionality (and security) to other developers. More recently, this sentiment about API developers was echoed by Green and Smith [9], who emphasised that developers typically focus on functionality and expect APIs to be secure by default. One issue

<sup>1</sup><https://www.bls.gov/ooh/computer-and-information-technology/software-developers.htm>

with this argument is that it treats software developers as a homogeneous population with little security awareness, but expects API developers to be somehow more security conscious. API developers are as human as software developers, subsequently they are susceptible to the same cognitive and social biases [4, 16]. Rather than assigning responsibility to different groups, we should identify the psychological dimensions associated with good security behaviours and seek to promote these skills in software learning development.

### 1.1 Motivation and Rationale

The primary motivation is to understand how soft skills relate to software development, and how people’s skills develop and exhibit in software development environments. Of particular interest are the behavioural changes exhibited by relative novices during skill development, compared to more experienced developers. What potentially incorrect, but intuitive actions are ultimately suppressed through experience? What habits are built and how do these originate? What soft skills are required for secure software development and how do these skills manifest and evolve?

In recent years, increased attention has turned towards the psychology of software developers, particularly in relation to security [17]. Security vulnerabilities typically leverage psychological processes [21], via cognitive processes (such as exploiting expected use cases), or through exploiting heuristic use. If adversaries exploit developers’ behaviour, it is important to identify the soft skills involved and find ways in which these behaviours can be changed through psychological interventions, which can be taught to novice and experienced developers alike.

The project is motivated to provide practical impact through developing teaching materials for Computer Science courses. Incorporating psychological interventions into pedagogy will allow for the development of soft skills and security conscious behaviours in future generations of software developers.

### 1.2 Contribution

The main aim of this research is to understand how the software learning development process occurs and how behaviours change and evolve. By identifying these processes, we are better placed to encourage positive behavioural changes, resulting in more efficient code development. The project also seeks to investigate key soft skills that affect secure code production. As a result, psychological interventions can be developed for promoting better security practices. To practically encourage relevant behavioural changes in early-stage software developers (e.g. Computer Science students), pedagogical materials will be developed for education, with the aim for these to be incorporated into teaching practices.

## 2 RESEARCH QUESTIONS

RQ1: What soft skills are considered important for computing and security practices?

RQ2: How do soft skills evolve and develop in novice software developers with time and experience?

RQ3: How do software development behaviours evolve and change with experience?

RQ4: How can relevant soft skills be incorporated into pedagogical practice to promote security behaviours?

## 3 WORK PLAN

The research incorporates a breadth of analysis methods, including qualitative and quantitative approaches. This methodological pluralism allows for a broad range of information to be drawn from the data that would otherwise not be possible with a restricted methodology. The doctoral research will look at data collected through interviews, surveys, data scraping and behavioural studies. Analysis will be varied and include statistical modelling, natural language processing modelling, and qualitative approaches, such as thematic analysis. Not only will it provide broader interpretation to findings within this space, it allows for stronger links to other work in similar research spaces.

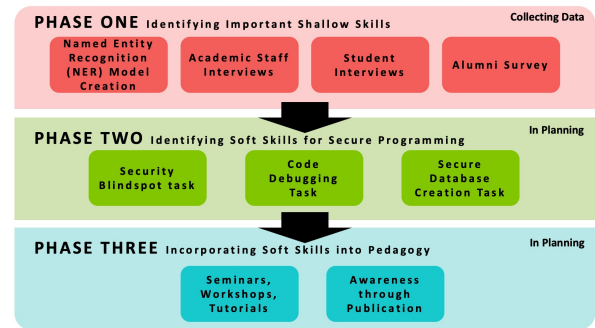


Figure 1: Outline of research phases within the planned work.

The research can be split into three phases, as illustrated by Figure 1. The first phase identifies shallow skills as taught in Computer Science undergraduate courses. It seeks to understand how shallow skills are perceived by current students, staff and alumni.

The second phase will build directly on the first phase. By identifying soft skills linked to shallow skills through previous research, phase two aims to build relationships between secure coding behaviours and soft skills. This will be achieved through empirical, lab-based research involving manipulation and measuring of soft skills and programming tasks.

Finally, the third phase will focus on incorporating findings on soft skills into pedagogical materials. This phase will measure the effectiveness of teaching these ideas, with the aim to raise awareness and increase the understanding of the psychological traits of software developers.

A stand-alone study is also being conducted into cognitive reflection and risk perception in software developers and computer scientists, see section 3.4. This will fit in with the phase two work.

Following open science practices, the research will include pre-registrations, data sharing and reproducible analysis scripts. This will be managed through the use of the Open Science Framework<sup>2</sup> and provision of Docker containers with reproducible workflows.

### 3.1 Phase One: Identification of Shallow Skills

The first phase is currently in the data collection stage. This phase is comprised of four research projects: an examination of core modules

<sup>2</sup>www.osf.io

in computer science programmes as taken from university websites; academic staff interviews on how they view shallow skills being taught; longitudinal interviews with current students on how they develop their shallow skills over an academic year; and an alumni survey of computer science graduates and psychology graduates, collecting data on their perceived importance of shallow skills.

**3.1.1 Curriculum Examination.** The online course information and core module descriptions for Computer Science and Psychology undergraduate courses were collected from eight UK universities belonging to the N8 research group (Durham, Lancaster, Leeds, Liverpool, Manchester, Newcastle, Sheffield and York).

To identify shallow skills in natural language texts, a named entity recognition (NER) model will be developed. Similar work has been created [6], but without the granularity attempted here. In efforts to further understand covariance of shallow skills and language used around them, the NER model weights can be analysed further, including factor analysis to find highly correlated skills. A preregistration, providing details on the data collection can be found at <https://osf.io/qcw3n>.

**3.1.2 Alumni Survey.** Lancaster University undergraduate alumni from Computer Science and Psychology were contacted to take part in a survey. Participants were asked to rank shallow skills for their importance in current employment. A Psychology sample were used as a comparative group, particularly when considering the less vocational nature of psychology undergraduate degrees (inferred from software development roles attained with a minimum education of a bachelor's degree<sup>3</sup>, compared to a minimum education of a postgraduate degree for most psychology roles<sup>4</sup>).


Data analysis will focus on loglinear models, correspondence analysis and exploratory factor analysis to identify the key shallow skills for computer science graduates compared to psychology graduates. The preregistration is found at <https://osf.io/5qb6a>.

**3.1.3 Interviews with Staff and Students.** These two projects are planned, and interview schedules will be arranged for times when teaching volume is low for staff, and longitudinal student surveys will begin in line with the start of an academic year.

Staff interviews will look to identify soft skills considered important by teaching staff and how these are conveyed to students in teaching materials. Student interviews will be conducted over the course of the academic year, following the same students to identify the way in which they recognise and develop shallow skills. Analysis for both interview studies will use thematic and content analysis to extract relevant information.

**3.1.4 Data Analysis of Phase One.** The results from the individual projects in phase one can be cross-examined to identify areas of shallow skills that are of most interest. The combined data can be used to understand the transmission of ideas from academic staff to students to what they take into the workplace, (see figure 2). Understanding the development of these skills and their importance can be used in the second phase. Analysis of data is in planning.

*Goal: to identify the shallow skills considered as important within the transmission of skills in the pedagogical process. Measured through various quantitative and qualitative methods.*



**Figure 2: Transmission of shallow skills can be expected to develop whereby teaching staffs' understanding of which skills are important feed into the skills students pick up on, which are reflected in alumni use of these skills.**

## 3.2 Phase Two: Behavioural Studies

The second phase is in early design stages, but will focus on empirical behavioural research, based on the findings from phase one. The exact soft skills to be included is dependent on phase one findings, as it is important to focus on the skills that are most likely to have the biggest effect on coding behaviours.

One study will investigate API blindspots, which can be defined as a misunderstanding or misrepresentation of API function security, resulting in vulnerabilities [16]. Using Python snippets from Brun et al (2021) [4], and measuring soft skills through cognitive tasks (e.g. the cognitive reflection test [7]), relationships can be drawn between soft skills and API blindspot awareness.

Similar studies, using different programming paradigms (such as code debugging, or secure password database creation) will also be used. It is important to understand the stability of soft skills across a range of development and security-related tasks.

**3.2.1 Data Analysis of Phase Two.** Data analysis for phase two experiments will be predominantly quantitative, using mixed effects models for group comparisons. These can be used to measure relationships between soft skills and security behaviours. Preregistrations will be published in due course.

*Goal: to identify and measure the effect of soft skills on security behaviours. Measured through mixed effect modelling and group comparisons.*

## 3.3 Phase Three: Inclusion in Pedagogy

For the final phase, the focus will be on the development of pedagogical materials for introducing students to the soft skills necessary for secure programming. This phase has not yet reached planning, as it relies on the work of phase two to be near completion. This will be achieved through seminars or workshops as methods to introduce the soft skills, to encourage students to engage with the psychology behind software development. Effectiveness of sessions will likely be measured through participant feedback.

To further disseminate research findings and promote inclusion of soft skills into current pedagogical materials, engagement

<sup>3</sup><https://nationalcareers.service.gov.uk/job-profiles/software-developer>

<sup>4</sup><https://nationalcareers.service.gov.uk/job-profiles/psychologist>

through publication will be pursued. By raising awareness of research through publication, conferences and posters, along with the provision of basic materials for others to work with, phase three looks to create a meaningful impact in the domain of software learning development.

*Goal: to develop and deliver teaching materials in order to promote soft skills within computer science curricula. Measured through student engagement and feedback.*

### 3.4 Risk Perception and Cognitive Reflection

In this individual differences study aligned with phase two research, groups of professional software developers and computer science students were compared regarding risk perception in software. Participants completed a cognitive reflection test, a risk-oriented decision task, and answered qualitative questions about how they understand risk in software development.

Cognitive reflection is a person's ability to inhibit intuitive responses in favour of more reflective responses, indicating their skill in reflective thinking in search of a correct answer. Cognitive reflection was measured through the Cognitive Reflection Test (CRT) [7]. This is a three question test, including items such as, "A paperclip and an elastic band cost £1.10 in total. The elastic band costs £1 more than the paperclip. How much does the paperclip cost?" The intuitive answer is 10 pence, but upon reflection the correct answer is 5 pence. The risk-orientation task focussed on how participants view susceptibility of themselves and the "average developer" when considering security vulnerabilities as listed by OWASP (e.g. SQL injection). Data analysis is in progress. The preregistration can be found at: <https://osf.io/zbqe4>.

*3.4.1 Data Analysis of Risk Perception Study.* Data will be analysed through quantitative measures, such as linear modelling, along with more qualitative methods, including thematic and content analyses.

*Goal: to identify potential relationships between risk-related behaviours in software development and cognitive reflection. Measured through linear modelling and group comparisons.*

### 3.5 Validity Threats and Controls

Validity threats to the research are broadly discussed, relevant to the project overall. More granular considerations are included in preregistration documents.

One key threat is the consideration of software developers as a population. It is easy to treat developers as a homogeneous population who demonstrate similar characteristics, subsequently making approaches to promoting security behaviours intolerant to variance within the population. This can be controlled through mixed effect models, where population characteristics can be included in the analysis to identify the effect these have on behaviours.

This is a relatively new research field [17], so much of the planned research is exploratory. This can often lead to a series of analysis methods being used, increasing type I errors. To control for this, preregistration procedures are published prior to data collection. Open data and reproducible analysis scripts will be uploaded following study completion, to allow replication and to confirm findings.

Using a range of methodologies, as highlighted in the work plan, may result in a trade-off between breadth of analysis and depth of analysis. To control for this, analysis plans and research choices will

be well considered through the use of preregistration documents. By considering methodologies prior to execution, the connection between the studies within the wider research can be well justified.

Another threat to validity is the generalisation to different programming languages or work cultures. Not all languages have similar structures, and differences have been shown in security awareness between Java and Python APIs [4]. This can be controlled for by acknowledging that results may only apply to a single language. By focussing on Python, which is the most popular language<sup>5</sup>, findings will have relevance to many developers. The inclusion of preregistration documents, materials and analyses will also allow for replications, either directly or conceptually.

The tasks used in the second phase are designed to provide consistency across participants, reducing task variance and improving statistical power. This comes at a cost, which is that the tasks are less industry-specific, reducing the validity. This PhD research is specifically focussed on the software learning process, and work beyond the PhD may look into more industry specific tasks, or applying similar research to different programming languages.

## 4 RELEVANT PRIOR WORK

In this section the current literature relevant to the research is discussed. This is not an exhaustive literature review, but aims to identify key research influencing the doctoral research.

### 4.1 Phase One

In phase one, key research identified shallow skills in software development, such as Matturro et al. (2019) [14], who conducted a systematic analysis and identified 23 separate skills. Similarly, Stevens and Norman (2016) looked at job adverts to identify the most important shallow skills for developers [20]. These research papers provided context for the important shallow skills.

Groeneveld et al. [10] analysed computer science curricula for modules that taught shallow skills explicitly, but did not look into the implicitly taught skills in all modules. This motivated the investigation of the course curricula for text relevant to shallow skills.

Finding an absence of research that provided associations between shallow skills and soft skills is the motivation for phase one. The literature search for phase one has found little evidence of work associating security awareness and soft skills.

### 4.2 Phase Two

In the second phase, a series of work has been carried out concerning API blindspots and developers' use of heuristics when evaluating software code. Oliveira et al. (2018) [16] highlighted this issue with Java puzzles, finding that security blindspots in code snippets were difficult to identify, possibly due to developers' expectation of APIs being secure as default. Brun et al. (2021) [4] followed this work with a replication using Python code. They found that developers who exhibited better long term memory recall were more successful in solving puzzles with blindspots. They found that short term memory, memory span and episodic memory had no effect on solving the puzzles. Other works that touch on psychology in security include Hallett et al (2021) [12], where boosting security awareness through requiring planning promoted a small effect

<sup>5</sup><https://www.tiobe.com/tiobe-index/python/>

on security, and Shreeve et al (2020) [19] who identified decision making processes related to cybersecurity.

### 4.3 Phase Three

For the third phase, Taylor-Jackson et al. (2020) [21] advocated including psychology in security education, particularly when considering that vulnerabilities are often psychological in nature (e.g. phishing, API blindspots). They discuss the benefits of exposing computer scientists to the different ideas and styles of thinking found within psychology. There are also wider calls for inclusion of soft skills in university education [11]. It is important that the findings from the first two research phases are used for positive impact and one immediate way to achieve this, is to answer the calls for increasing soft skill teachings in cybersecurity courses to benefit future software developers.

### 4.4 Risk Perception

For this study, key items are papers on cognitive reflection by Frederick (2005) [7] and Thomson and Oppenheimer (2016) [22]. Combined with the understanding that developers are often not the most security conscious, as highlighted by Acar et al (2017) [1], it is clear that the understanding of risk by developers in a software context is poorly understood in relation to cognitive measures.

## 5 CURRENT STATUS

### 5.1 Early Results Analysis

Some of the preliminary results from the risk perception study (section 3.4) are mentioned here. The third hypothesis stated in the preregistration is examined here, "Mean scores closer to zero on the novel OWASP risk task will be found with higher scores of cognitive reflection". Data from 143 (70 students, 73 developers) participants is used.

The OWASP task is a measure devised for this study where participants were asked to respond to two sets of questions, the first asking about the percentage of web applications they believe to be created by others that suffer from one of the top five OWASP vulnerabilities (injection flaws, broken authentication, sensitive data exposure, XML External Entity and broken access control). Then following a separation task, participants were then asked to rate the percentage of web applications that they had developed that suffered from the same vulnerabilities. Scores closer to 100 indicate high optimism that they do not produce flawed products, scores near 0 indicate similar levels of flaws in both their own and other people's products, and scores approaching -100 indicate beliefs that their own work is highly susceptible to these vulnerabilities.

To test the hypothesis above, a linear regression was run to see whether CRT scores significantly predicted OWASP vulnerability scores and whether this differed between the two populations. The model formula was "Vulnerability ~ CRT". The overall regression was statistically significant ( $R^2 = .05$ ,  $F(3, 139) = 15.13$ ,  $p = .017$ ). Estimates, values and significance of model items can be seen in Table 1. Despite significant terms in the model summary, the variance explained by the model is negligible (~5%) and further models will need to be developed to explain more variance in these scores.

Figure 3 shows the distribution of vulnerability scores for each level of CRT score as a box plot. Post-hoc Tukey tests identified no

**Table 1: Coefficients, t-values and p-values for the linear regression of CRT predicting OWASP vulnerability**


	Estimate	Std. Error	t value	$p^*$
Intercept	11.00	3.13	3.52	<.001***
CRTscore1	14.08	4.59	3.07	.003**
CRTscore2	5.61	4.86	1.15	.250
CRTscore3	10.67	4.91	2.17	.031*

\*Significant alpha values of <.001 indicated by \*\*\*

significant differences between any of the groups with all  $p > .05$ , except for those who scored zero and those who correctly scored one, adjusted  $p = .014$ . This indicates that there is little significant relationship between CRT scores and results on the novel OWASP risk task.

What is noted with the Vulnerability scores, and can be seen in Figure 3, is that most scores on the OWASP task, regardless of CRT scores, are around or above 0. One-way t-tests on the Vulnerability scores were run for both the developer and the student samples. In the Developer sample (mean score = 17.48), the scores were significantly higher than 0,  $t(69) = 7.16$ ,  $p < .001$ . Similarly for the student sample (mean score = 18.94), scores were significantly higher than 0,  $t(69) = 7.26$ ,  $p < .001$ .

This finding is indicative of optimism bias [18], suggesting that both professionals and student developers consider themselves to be better than average at preventing these OWASP-listed security issues. A score of zero would indicate respondents understand they were average, but higher scores suggest an over-optimistic outlook on their own abilities, which could lead to a more relaxed view on these security issues. These findings will be further developed and discussed in future publications.



**Figure 3: Box plot of mean OWASP vulnerability scores by CRT score split by population.**

## 5.2 Next Steps

In the short term, the next steps are to continue with phase one data collection, and planning of phase two. It is intended that the research will progress according to the research phases outlined above. Following the completion of the doctoral work, future work would include the investigation of psychological interventions for developing soft skills and measuring their impact on security behaviours in longitudinal research.

The steps beyond the PhD research as outlined above is to focus on the skills that explain the largest variance in secure coding behaviours, and seek to identify the best ways to promote continued, stable use of these behaviours as opposed to short-term changes (such as those achieved through nudging, e.g. IDE pop-ups serving as reminders to look for blindspots).

## 6 CONCLUSION

This paper provides an overview of the planned work within the PhD research titled "The Soft Skills of Software Learning Development: the Psychological Dimensions of Computing and Security Behaviours". The research is diverse in both aims and processes, ranging from thematic analysis of interviews, to modelling relationships between psychological dimensions and security issues, to incorporating the findings into pedagogy.

This project seeks to investigate the software learning development process; to better understand the behavioural changes and soft skill development of both computing and security behaviours. By identifying these changes, and when and how they develop, we can seek to promote these changes earlier in the learning cycle, allowing for more effective learning and encouraging positive behaviours for software development. In doing so, not only can we exhibit greater awareness of the psychology behind secure software development, we can develop interventions for encouraging these secure behaviours, reducing the likelihood of these security vulnerabilities. All public preregistrations, published data, analyses, and links to further research outputs will be accessible from <https://osf.io/v93zt>.

## REFERENCES

- [1] Yasemin Acar, Christian Stransky, Dominik Wermke, Charles Weir, Michelle L. Mazurek, and Sascha Fahl. 2017. Developers Need Support, Too: A Survey of Security Advice for Software Developers. In *2017 IEEE Cybersecurity Development (SecDev)*. IEEE, Cambridge, MA, USA, 22–26. <https://doi.org/10.1109/SecDev.2017.17>
- [2] Anne Adams and Martina Angela Sasse. 1999. Users Are Not the Enemy. *Commun. ACM* 42, 12 (Dec. 1999), 40–46. <https://doi.org/10.1145/322796.322806>
- [3] Faheem Ahmed, Luiz Fernando Capretz, Salah Bouktif, and Piers Campbell. 2015. Soft Skills and Software Development: A Reflection from the Software Industry. *International Journal of Information Processing and Management* 4, 3 (July 2015), 171–191. <https://doi.org/10.4156/ijipm.vol14.issue3.17> arXiv:1507.06873
- [4] Yuriy Brun, Tian Lin, Jessie Elise Somerville, Elisha Myers, and Natalie C. Ebner. 2021. Blindspots in Python and Java APIs Result in Vulnerable Code. *arXiv:2103.06091 [cs]* (March 2021). arXiv:2103.06091 [cs]
- [5] Luiz Fernando Capretz and Faheem Ahmed. 2018. A Call to Promote Soft Skills in Software Engineering. *Psychology and Cognitive Sciences - Open Journal* 4, 1 (Aug. 2018), e1–e3. <https://doi.org/10.17140/PCSOJ-4-e011> arXiv:1901.01819
- [6] Silvia Fareri, Nicola Melluso, Filippo Chiarello, and Gualtiero Fantoni. 2021. SkillNER: Mining and Mapping Soft Skills from Any Text. *Expert Systems with Applications* 184 (Dec. 2021), 115544. <https://doi.org/10.1016/j.eswa.2021.115544> arXiv:2101.11431
- [7] Shane Frederick. 2005. Cognitive Reflection and Decision Making. *Journal of Economic perspectives* 19, 4 (2005), 25–42.
- [8] Steven Furnell and Matt Bishop. 2020. Addressing Cyber Security Skills: The Spectrum, Not the Silo. *Computer Fraud & Security* 2020, 2 (Feb. 2020), 6–11. [https://doi.org/10.1016/S1361-3723\(20\)30017-8](https://doi.org/10.1016/S1361-3723(20)30017-8)
- [9] Matthew Green and Matthew Smith. 2016. Developers Are Not the Enemy!: The Need for Usable Security APIs. *IEEE Security Privacy* 14, 5 (Sept. 2016), 40–46. <https://doi.org/10.1109/MSP.2016.111>
- [10] Wouter Groeneveld, Brett A. Becker, and Joost Vennekens. 2020. Soft Skills: What Do Computing Program Syllabi Reveal About Non-Technical Expectations of Undergraduate Students?. In *Proceedings of the 2020 ACM Conference on Innovation and Technology in Computer Science Education (ITICSE '20)*. Association for Computing Machinery, New York, NY, USA, 287–293. <https://doi.org/10.1145/3341525.3387396>
- [11] Sandra Patricia Guerra-Báez. 2019. A Panoramic Review of Soft Skills Training in University Students. *Psicologia Escolar e Educacional* 23 (2019), 1–10. <https://doi.org/10.1590/2175-35392019016464>
- [12] Joseph Hallett, Nikhil Patnaik, Benjamin Shreeve, and Awais Rashid. 2021. "Do This! Do That!, And Nothing Will Happen" Do Specifications Lead to Securely Stored Passwords?. In *Proceedings of the 43rd International Conference on Software Engineering (ICSE '21)*. IEEE, Madrid, Spain, 486–498. <https://doi.org/10.1109/ICSE43902.2021.00053>
- [13] Janet Liebenberg, Magda Huisman, and Elsa Mentz. 2014. Knowledge and Skills Requirements for Software Developer Students. *International Journal of Social, Behavioral, Educational, Economic, Business and Industrial Engineering* 8, 8 (2014), 6.
- [14] Gerardo Maturro, Florencia Raschetti, and Carina Fontán. 2019. A Systematic Mapping Study on Soft Skills in Software Engineering. *Journal of Universal Computer Science* 25, 1 (2019), 26.
- [15] João Eduardo Montandon, Cristiano Politowski, Luciana Lourdes Silva, Marco Tulio Valente, Fabio Petrillo, and Yann-Gaël Guéhéneuc. 2021. What Skills Do IT Companies Look for in New Developers? A Study with Stack Overflow Jobs. *Information and Software Technology* 129 (Jan. 2021), 106429. <https://doi.org/10.1016/j.infsof.2020.106429>
- [16] Daniela Seabra Oliveira, Tian Lin, Muhammad Sajidur Rahman, Rad Akefirad, Donovan Ellis, Eliany Perez, Rahul Bobhate, Lois A DeLong, Justin Capps, and Yuriy Brun. 2018. {API} Blindspots: Why Experienced Developers Write Vulnerable Code. In *Fourteenth Symposium on Usable Privacy and Security (SOUPS) 2018*. USENIX Association, Baltimore, MD, USA, 315–328.
- [17] Irum Rauf, Marian Petre, Thein Tun, Tamara Lopez, Paul Lunn, Dirk Van Der Linden, John Towse, Helen Sharp, Mark Levine, Awais Rashid, and Bashar Nuseibeh. 2021. The Case for Adaptive Security Interventions. *ACM Transactions on Software Engineering and Methodology* 31, 1 (Sept. 2021), 9:1–9:52. <https://doi.org/10.1145/3471930>
- [18] Tali Sharot. 2011. The Optimism Bias. *Current Biology* 21, 23 (Dec. 2011), R941–R945. <https://doi.org/10.1016/j.cub.2011.10.030>
- [19] Benjamin Shreeve, Joseph Hallett, Matthew Edwards, Pauline Anthonysamy, Sylvain Frey, and Awais Rashid. 2020. "So If Mr Blue Head Here Clicks the Link." Risk Thinking in Cyber Security Decision Making. *ACM Transactions on Privacy and Security* 24, 1 (Nov. 2020), 5:1–5:29. <https://doi.org/10.1145/3419101>
- [20] Matt Stevens and Richard Norman. 2016. Industry Expectations of Soft Skills in IT Graduates: A Regional Survey. In *Proceedings of the Australasian Computer Science Week Multiconference (ACSW '16)*. Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/2843043.2843068>
- [21] Jacqui Taylor-Jackson, John McAlaney, Jeffrey L. Foster, Abubakar Bello, Alana Maurushat, and John Dale. 2020. Incorporating Psychology into Cyber Security Education: A Pedagogical Approach. In *Financial Cryptography and Data Security (Lecture Notes in Computer Science)*, Matthew Bernhard, Andrea Bracciali, L. Jean Camp, Shin'ichiro Matsuo, Alana Maurushat, Peter B. Rønne, and Massimiliano Sala (Eds.). Springer International Publishing, Cham, 207–217. [https://doi.org/10.1007/978-3-030-54455-3\\_15](https://doi.org/10.1007/978-3-030-54455-3_15)
- [22] Keela S Thomson and Daniel M Oppenheimer. 2016. Investigating an Alternate Form of the Cognitive Reflection Test. *Judgment and Decision making* 11, 1 (2016), 99.
- [23] John Towse, Mark Levine, Marian Petre, Arosha Bandara, Tamara Lopez, Awais Rashid, Irum Rauf, Helen Sharp, Thein Tun, Dirk van der Linden, and Bashar Nuseibeh. 2020 (in press). The Case for Understanding Secure Coding as a Psychological Enterprise. *Cyberpsychology, Behavior, and Social Networking* (2020 (in press)).
- [24] Glenn Wurster and Paul van Oorschot. 2008. The Developer Is the Enemy. In *Proceedings of the 2008 New Security Paradigms Workshop (NSPW '08)*. Association for Computing Machinery, New York, NY, USA, 89–97. <https://doi.org/10.1145/1595676.1595691>