"Do It to Know It": Reshaping the Privacy Mindset of Computer Science Undergraduates

Maisha Boteju
University of Auckland
New Zealand
mbot450@aucklanduni.ac.nz

Danielle Lottridge
University of Auckland
New Zealand
d.lottridge@auckland.ac.nz

Thilina Ranbaduge
CSIRO's Data61
Canberra, Australia
thilina.ranbaduge@data61.csiro.au

Dinusha Vatsalan
Macquarie University
Australia
dinusha.vatsalan@mq.edu.au

Ni Ding
University of Auckland
New Zealand
ni.ding@auckland.ac.nz

Abstract

Software applications, while being an integral part of the modern world, pose significant threats to end-user privacy. Thus, computer professionals require knowledge and skills to develop privacy-aware software. However, undergraduate computing degree programs often lack privacy-focused curricula that can cultivate this ability in the future workforce. Therefore, we designed a privacy curriculum informed by the common challenges that computing professionals face when developing privacy-embedded software. It guides students in realising the need for privacy, identifying privacy protection mechanisms and programming Privacy Enhancing Technologies (PETs). We piloted the curriculum for third-year Computer Science undergraduates at the University of Auckland, New Zealand. The curriculum was evaluated using course assessments and surveys conducted before and after the lessons. Overall, the students improved their understanding of privacy, especially technical aspects. Most of them valued the applied learning experience of the programming lessons yet showed distinct views on task completion difficulty and motivation to do programming. Students recognised that privacy should be integral to their skill set by confirming the importance and relevance of the lessons. However, their perceived responsibility in privacy protection varied depending on their intention to take proactive measures. Based on the results, the paper suggests improvements to the proposed curriculum.

Keywords

privacy, Privacy Enhancing Technologies, Computer Science education, pedagogy, software development, data protection

1 Introduction

The vast accumulation of data through software applications expands the attack surface on people's privacy. As such, researchers have explored ways to enhance laypeople's privacy literacy to promote more informed data-sharing practices [4, 6, 41, 64]. The end-user perspective on privacy has limitations, as users are often unaware of the complexities involved in data handling behind software systems [1, 23]. Given this context, it is vital to develop software with privacy constraints embedded as a core feature. This approach is also emphasised by data protection regulations (e.g., General Data Protection Regulation [19]), principles of global organisations (e.g., the United Nations [21]) and standardisation bodies (e.g., International Organisation for Standardisation [28]).

Computer Science (CS) professionals face challenges when integrating privacy into software. They struggle to define privacy and differentiate it from security [27]. Many are unaware of different privacy protection mechanisms, while some struggle to select and apply suitable ones [47, 59, 61]. Some lack the mindset to empathise with users or realise the consequences of privacy violations [2, 18, 39, 53]. Due to these challenges, privacy is not a primary concern in CS professionals' daily workflow, and they often transfer the responsibility of privacy preservation to upper management or privacy experts [2, 3, 29, 68]. Consequently, privacy remains unaddressed or an afterthought in software development [29]. Thus, Computer Science (CS) professionals, not limited to privacy specialists, require knowledge and skills to address data privacy [33, 40].

A well-structured privacy curriculum at the undergraduate level can train future professionals to address the above challenges [27]. To date, institutions and research efforts have given limited attention to improving the privacy education of CS undergraduates [46]. Emerging technical privacy curricula and teaching methods are not readily available. While existing efforts are valuable, they fail to align the proposed lessons or teaching methods with the industry demands (e.g., threat modelling, implement privacy protection techniques) and lack in-depth evaluation.

This paper introduces a novel curriculum on privacy that teaches CS undergraduates the importance of privacy and the technical aspects of privacy protection. The lessons guide the students through **1. introduction to privacy** (definition, societal aspect, regulations, contextual privacy [50]), **2. data privacy protection** (digital footprint, Privacy by Design [11], LINDDUN threat modelling [76]), **3. PETs** (theoretical background, known weaknesses, secondary benefits, selecting suitable PETs) and **4. programming PETs** (pseudonymisation, k-anonymity and differential privacy). These lessons were designed to mitigate the difficulties computing professionals face when addressing privacy concerns. Lesson 1 aligns with the challenges: define privacy, differentiate it from

Proceedings on Privacy Enhancing Technologies 2025(4), 969–995

https://doi.org/10.56553/popets-2025-0167

security, and understand the consequences of privacy breaches. Lessons 2 and 3 teach foundational privacy protection mechanisms. Lesson 4 aligns with the challenge of applying privacy protection mechanisms. As the lessons start from conceptual levels and span towards practical levels, they offer students a chance to improve their knowledge progressively.

We piloted the curriculum for third-year CS undergraduates at the University of Auckland, New Zealand. The lessons were included in the COMPSCI 316 cybersecurity course and taught over two weeks. We aimed to answer the following research questions:

RQ1 : How has the proposed curriculum improved the privacy knowledge of Computer Science undergraduates?
RQ2 : How do Computer Science undergraduates perceive privacy programming lessons?
RQ3 : How can the proposed curriculum help Computer Science undergraduates realise their role in privacy protection?

To investigate the research questions, we evaluated the course using survey data (before and after the lessons) and assessments (optional exercises, group assignment, mid-semester exam, and final exam). The lesson slides, assessments, and programming exercises are available in the repository at https://github.com/MaishaB/undergraduate-privacy-curriculum.

The curriculum improved the students' understanding of privacy. There was a significant improvement in their ability to identify privacy threats and their perceived confidence and skill in implementing privacy protection. This outcome is notable, given their limited technical knowledge before the course. However, students needed more understanding of offline privacy threats and the distinction between security and privacy. We also identified five distinct learner groups based on the perceptions towards PETs programming: *capable but unengaged, moderately engaged, hands-on, perseverant, and struggling*. This differentiation will aid lesson customisation in future. The students found the content relevant and important to CS undergraduates. While most students felt a sense of responsibility towards privacy protection, their intention to take proactive action differed. Based on these findings, we proposed improvements to the content, structure and execution of the curriculum.

2 Related Work

To identify research contributions towards privacy education in undergraduate CS curricula, we searched in the ACM digital library, IEEE Explorer, Web of Science, Springer Link, SAGE journals and Google Scholar. We used the search string *"privacy" AND ("education" OR "teach" OR "curriculum") AND ("computer science" OR "software engineering" OR "information technology" OR "computing")*. After screening the search results of the last decade, we retrieved five publications; two focused on proposing curricula [46, 57] and three examined teaching methods (e.g., lab exercises) [38, 62, 63].

Moore et al. [46] proposed a privacy course covering ethical, technical, and legal aspects. It included adversarial thinking as part of the exercises, where students implemented server-side tracking techniques. The authors also included exercises to code k-anonymity, l-diversity and differential privacy, providing hands-on experience with privacy protection. The curriculum improved the students' confidence in designing privacy-preserving systems. Redd et al. [57] integrated sociotechnical privacy concepts into a capstone project course covering regulations, privacy settings, and threat analysis. Students exhibited a notable improvement in privacy awareness compared to those in traditional capstone courses.

Few studies explored creative teaching methods to enhance student engagement with privacy concepts. Shilton et al. [63] suggested a role-playing activity leveraging privacy by design practices. Students worked as a group to decide policy and technical changes when a mobile application shifts its market platform (e.g., Android, iOS) to a more restrictive one. The intervention was well received and sparked student interest in learning more about ethics in technical work.

Li et al. [38] introduced a lab-based learning approach based on location privacy. This labware provided a tool that students could use to observe how location anonymisation impacts the accuracy of location-based services (LBS). The lab was perceived as effective and increased students' awareness, interest, and understanding of privacy disclosure, LBS, and anonymisation. Further, Rydal et al. [62] used experiential learning in which students collected, visualised, and reflected on their movement data. This method fostered awareness of privacy considerations and cultivated a sense of responsibility and caring towards protecting data.

Additionally, we manually reviewed the course catalogues and artifacts of the top 50 universities in the 2025 Times Higher Education rankings[1] for the computer science field (Appendix E, Table 8). Only 14% of these universities offered dedicated privacy courses, while 32% integrated privacy with other subjects (e.g., ethics, security). This limited focus on privacy education is consistent with observations in [46]. Most courses focused on legal, theoretical, or policy aspects of privacy, covering topics like social impact, regulations, and privacy risks. Practical components, like threat modelling and privacy protection implementation, were largely absent.

Related research [38, 46, 57, 62, 63] and tertiary curricula lay an important foundation in privacy education. In comparison, Table 8 highlights the following key contributions of our approach:

Pedagogical: Our curriculum was built to address industry challenges: inability to prioritise privacy [2, 18, 39, 53], lack of technical skills [47, 59, 61], and reluctance to take responsibility for privacy concerns [2, 3, 29, 68], aligning it with professional needs. Compared to others, our curriculum is progressive, including lessons and assessments moving from foundational concepts to practical skills. It stands out for balancing both user and professional perspectives of privacy, an aspect overlooked in related efforts. It expands the technical breadth by integrating threat modelling, a wide range of PETs, and hands-on programming, a combination absent in related work. Finally, we made our lessons and assessments open-access to support adoption, adaption and research.

Scientific: We contribute to an underexplored area, 'undergraduate privacy education', by introducing a curriculum with evidence-based design and the most detailed evaluation to date. As shown in Table 8 in Appendix E, we stand out through our use of diverse evaluation methods and multifaceted findings. Generally, related research [38, 46, 57, 62, 63] reports pre-post knowledge differences

[1] https://www.timeshighereducation.com/world-university-rankings/latest/world-ranking

using quantitative or qualitative methods. The qualitative evaluation in [62] offers thematic insights but lacks measurable evidence. The quantitative evaluations show statistical improvement only across some intended learning goals [46, 57, 63], omit effect sizes [46], do not justify statistical test assumptions [46, 57], and use mean score changes limiting statistical validity [38]. Conversely, we achieve statistically significant improvement across all learning objectives and maintain credibility and validity by reporting effect sizes (medium to large), 95% confidence intervals, and validated test assumptions. Further, we capture students' perceptions of PETs programming and their perceived role in privacy protection, insights crucial for tailoring curricula that offer technical competency and a sense of professional involvement. Most importantly, we use these findings to interpret changes in students' privacy mindset, an important determinant of their future behaviour [13]. We also examine assessment quality and engagement, which are important for improving instruction but are missing in prior work. Moreover, contrary to the existing studies, we propose curriculum improvements based on our findings.

This methodological rigour sets our work apart and strengthens the credibility of our insights into curriculum effectiveness. While the findings may vary depending on the cohort or instructor, our evaluation strategy is generalisable due to its grounding in real-world privacy demands. Even CS courses outside the privacy domain can benefit from our approach (e.g., tailor content for different learners), further increasing the generalisability.

3 Methodology

The proposed privacy curriculum was piloted with third-year CS undergraduates at the University of Auckland, New Zealand. It was integrated into the existing COMPSCI 316 cybersecurity course (12 weeks). The privacy lessons were taught in weeks six and seven of the cybersecurity course using six hours of lectures and a tutorial session. Recorded lectures were provided to the students, allowing them to revisit the material. The tutorial session was a revision of the lectures, in which the students were given time to complete eight questions: six multiple choice and two essay type, followed by a discussion with a graduate teaching assistant. Prerequisites for the cybersecurity course and the lessons covered before week 6 is shown in Appendix B.

3.1 Learning Outcomes

We employed the backward design strategy of curriculum design [45], where the learning outcomes are first defined, followed by the teaching content creation. This method allowed us to align the lessons and assessments to the intended learning outcomes. To contextualise the learning outcomes of the curriculum, we used the privacy-related challenges computing professionals face, which were reported in the interview-based research studies [2, 3, 29, 47, 59, 61]. We defined three broader learning outcomes that provide knowledge and skills to counter these identified challenges: *1. Students know the need for privacy*, *2. Students know existing privacy protection mechanisms*, and *3. Students can implement privacy protection mechanisms*. "Students know the need for privacy" instils a sense of care in protecting others' privacy. Without this initial motivation, the students will not understand the importance of developing privacy-preserving software. "Students know existing privacy protection mechanisms" improves students' ability to identify which technique to use in different contexts when protecting privacy. Finally, "Students can implement privacy protection mechanisms", enhances their competency in integrating privacy measures into real-world applications.

The broader learning outcomes were further divided into more manageable and measurable ones (Table 1). They span different cognitive levels proposed in the revised Bloom's taxonomy [35]: remember, understand, apply, analyse, evaluate and create. For example, defining privacy is a remembering activity, while programming PETs covers applying, which is more advanced than remembering. Table 7 in Appendix A maps each learning outcome to its corresponding cognitive level. To create lessons and exercises that support these learning outcomes, we synthesised content from privacy education research [17, 38, 46, 57, 62, 63], other privacy research [8, 11, 15, 37, 50, 66, 76] and online resources [22, 49, 52, 54, 60]. Appendix A includes a detailed description of the lessons developed along these learning outcomes.

3.2 Data Collection

To answer the research questions, we collected data through surveys and assessments. The students were given two surveys: before (pre) and after (post) the privacy lectures. They were instructed to enter a pseudonym for their name during the pre-survey and post-survey so we could pair their responses anonymously. Regarding assessments, we used four ways to collect data to evaluate the learning outcomes, including optional exercises, mid-semester test (MST), assignment and the final exam. The IRB of the university approved the data collection and usage. All data were anonymised according to the ethics guidelines.

3.2.1 Pre-Survey. The pre-survey consisted of seven items (LP1 to LP7), including six Likert items and one open-ended question. The items corresponded with the learning outcomes. In the Likert items, the students had to self-rate themselves on different response scales (e.g., "I don't know" to "I'm an expert" or "not aware" to "very aware"). The 'neutral' option was removed from the responses as it was not conceptually meaningful for those questions, and students could have used it to avoid answering the item [32]. We used AhaSlides[2], an interactive presentation platform, to conduct the pre-survey. Students accessed it via a QR code, and responses were displayed after all submissions. By framing the activity as part of the learning experience, we aimed to promote more honest responses. The survey was provided as a Google form for the students who failed to attend the first lecture. To ensure the data quality, we removed the straightlining responses (giving the same response to consecutive questions [42]).

LP1 : What is "Privacy"?
LP2 : Rate your awareness regarding privacy threats in the offline world.
LP3 : Rate your awareness regarding privacy threats in the online world.

[2] https://ahaslides.com/

Table 1: Learning outcomes of the proposed curriculum and the methods used to evaluate those outcomes.

Broader Outcome		Learning Outcome	Evaluation			
			Optional Exercises	MST	Assignment	Final Exam
Students know the need for privacy	LO1	Students can define what privacy means	✓		✓	
	LO2	Students can explain why privacy is important for individuals, society, and organisations	✓		✓	
	LO3	Students can classify different privacy threats according to the LINDDUN threat model	✓	✓	✓	✓
	LO4	Students can generate a plan to improve the privacy climate of a software organisation			✓	
Students know existing privacy protection mechanisms	LO5	Students can distinguish privacy protection mechanisms from security protection mechanisms	✓		✓	
	LO6	Students can select suitable technical and non-technical privacy protection mechanisms for a given software application scenario	✓	✓	✓	✓
Students can implement privacy protection mechanisms	LO7	Students can implement pseudonymisation, k-anonymity and differential privacy using Python			✓	✓
	LO8	Students can justify their own decisions taken during the implementation of pseudonymisation, k-anonymity and differential privacy	✓		✓	✓

MST = Mid Semester Test

LP4 : Rate your ability to identify privacy vulnerabilities and privacy threats.
LP5 : Rate your confidence in trying to mitigate the identified privacy threats.
LP6 : Rate your skill level required to mitigate the identified privacy threats.
LP7 : Which mechanisms can be used to protect data privacy during software development?

3.2.2 Post-Survey. The post-survey included three main sections corresponding to the defined research questions: learning progress, perceptions of the programming lessons and perceptions of the role in privacy protection. In addition, it included an open-ended question at the end of the survey for additional comments. Since this survey was long, we avoided conducting it during the lectures and provided it through a Google Form. To motivate participation, we offered one bonus mark to each participant if the total number of responses exceeded 150. The survey was given during week 10, two weeks after the last privacy lecture, giving students time to process and reflect on the content.

Learning Progress. Students were first presented with a list of privacy topics covered in the course. They were asked to indicate whether they had learned about each topic before the course and, if so, to specify the learning resource: school, other university courses, external courses, social media, books, research papers, and other. We specifically added this question to the post-survey rather than the pre-survey, as students were expected to develop a clear understanding of the topics by that time. The topics included:

- The meaning of "Privacy"
- Social aspect of privacy
- Importance of privacy protection for individuals
- Importance of privacy protection for organisations
- Data privacy
- Privacy threats (LINDDUN or other)
- Different types of PETs
- Programming PETs

Then, to evaluate the progress along the learning outcomes, we included the same questions from the pre-survey (LP1 to LP7) in the post-survey. As a complementary measure, we asked the students to indicate the perceived improvement in their overall understanding of privacy using a seven-point scale.

LP8 : "The offered privacy lessons improved my understanding about privacy." On a scale of 1 to 7, how much do you agree with the above statement?

Perceptions of the Programming Lessons. In the post-survey, we included eight seven-point Likert items (PP1 to PP8) to retrieve students' feelings towards the programming lessons. To ensure the validity of the responses, we wanted to identify the students who attended (or watched) at least one programming lecture out of the three. For this, the survey asked students whether they attended (or watched) the lectures with options: No, Partially, or Yes.

PP1 : Too much work
PP2 : The exercises were hard
PP3 : It was interesting to convert theory into practice
PP4 : It helped me to understand more about the theory
PP5 : I wish I had more exercises
PP6 : It motivated me to program other PETs
PP7 : It showed me how privacy relates to software development
PP8 : It motivated me to try PETs in my software projects

Perceptions of the Role in Privacy Protection. In the final part, we wanted to investigate how students interpret their involvement in privacy protection. We included four questions: three seven-point Likert items (PR1 to PR3) and one open-ended question (PR4).

PR1 : How would you rate the relevance of the weeks 6 and 7 content to Computer Science undergraduates?
PR2 : How important do you think it is to include the weeks 6 and 7 content in undergraduate Computer Science courses?
PR3 : How well did the content from weeks 6 and 7 make you feel about your responsibility to protect privacy?
PR4 : How do you think the content covered in weeks 6 and 7 will help you in your future profession?

To validate the post-survey responses, we applied the following exclusion criteria.

EC1 : Straightlining
EC2 : Have learned PETs programming from school
EC3 : Responded 'No clue' to 'What is privacy?' but responded that they had learned privacy and programming PETs from various sources before the course
EC4 : Have learned different PETs and programming PETs from other sources before the course but failed to name a single PET as a privacy protection mechanism
EC5 : Have responded 'No' to 'Have you ever learned different PETs before the course?. However, responded 'Yes' to 'Have you ever learned programming PETs before the course?'

3.2.3 Assessments. To encourage students in discussions, we provided eight optional exercises and purposefully did not provide answers to some of the exercises. The MST (20 questions) and the final exam (50 questions) were multiple choice questions based tests. The number of privacy-related questions in the MST and the final exam were two and four, respectively. The assignment was released in the end of week six. It contained five questions: three essay type and two programming-related. Since addressing privacy concerns is considered a collaborative effort during software development [51], we wanted the students to build that experience by completing the assignment as a group. They were given one week to form groups of five, except for one group of six, due to the total of 221 students. We expected that the students' diverse societal backgrounds would lead to interesting brainstorming sessions. The students were asked to explain their individual contributions to help us understand their collaborative efforts.

3.3 Evaluation

We adopted a mixed methods approach combining quantitative and qualitative analyses to evaluate the curriculum.

3.3.1 Pre and Post Surveys. The survey responses were intended to be paired. However, only 25 students correctly entered matching pseudonyms in both surveys, limiting our ability to conduct evaluation methods that handle paired data.

Learning Progress. First, we cross-tabulated prior privacy knowledge and learning resources, which allowed us to identify knowledge gaps and assess the contribution of different resources to improve the privacy knowledge of CS students. Then, we analysed the learning progress using responses of LP1 to LP8.

LP1 to LP6 were Likert items intended for paired comparison. While the Wilcoxon signed-rank test [74] is suitable for evaluating such paired ordinal data, we did not apply it due to the limited number of paired responses. Thus, we analysed LP1 to LP6 responses using the Mann–Whitney U test [43]. This test was appropriate as the responses were unpaired, non-normally distributed, and ordinal [48]. The comparison of the response distributions is included in Appendix D, Figure 9. The null hypothesis for each Likert item was, "There is no significant difference between the distributions of the pre-survey and post-survey responses". The asymptotic significance (2-tailed p value) was used to validate the hypotheses. To quantify the strength of the difference between the distributions, we calculated eta-squared effect size [69] with 95% confidence intervals. The eta-squared value is interpreted as small (<0.05), medium (0.06 - 0.13) and large ($\geq$ 0.14) [30].

For the LP7 open-ended question, we visualised the responses through word clouds. Since the dataset sizes of pre- and post-surveys were different, our effort here was not to compare the frequencies of the answers. Instead, we observed interesting patterns in the responses, such as increased references to privacy-specific measures after the lessons.

Lastly, to reflect the improvement in students' understanding of privacy, we calculated the central tendency of item LP8 using the median measure.

Perceptions of the Programming Lessons. We performed Ward's hierarchical clustering [71] on Likert items PP1 to PP8 to identify different learner groups so that we can design more targeted programming lessons in the future. This clustering method is suitable for our analysis since we do not have a defined number of clusters before the clustering process [77]. We first reverse-scored the negatively worded Likert items. Then, to improve the interpretability of the clusters, we conducted dimensionality reduction using principal axis factoring [10] and identified three latent factors to represent the eight Likert items. Finally, hierarchical clustering was performed on these latent factors. The number of clusters was verified using the validation method presented in [77], in which the sudden spike in the agglomeration schedule coefficients is selected as the stage where clustering is stopped.

Perceptions of the Role in Privacy Protection. We evaluated the responses of PR1 to PR3 Likert items in two ways. First, to understand the perceived relevance (PR1) and importance (PR2) in learning privacy and perceived privacy protection responsibility (PR3), we calculated the central tendency of each item using the median. Second, to examine how these perceptions related to the overall understanding of privacy, we calculated the correlation between PR1, PR2, PR3 and LP8 (understanding of privacy). We used Kendall's tau-b correlation since the responses were ordinal, non-normal and displayed monotonic relationships [55]. The scatter plots showcasing the monotonic relationships are included in Appendix D, Figure 11. The strength of the correlation can be interpreted as negligible (0.0 - 0.05), weak (0.06 - 0.25), moderate (0.26 - 0.48), strong (0.49 - 0.7) and very strong ($\geq$ 0.71) [73].

Finally, we used Braun and Clarke's Reflexive Thematic Analysis (RTA) [9] to interpret the responses of PR4. To draw on subjective reflections, we used our experience in the software industry, privacy research, and teaching. We read the responses in multiple directions, top-down, bottom-up, and middle-out, to familiarise ourselves with the data and identify patterns without being influenced by a fixed reading order. When reading the responses, we maintained notes of our thoughts, which were later used to interpret the results. We then coded the data in two rounds. The first author independently

generated the initial codes, and then all authors refined the codes, replacing some with existing codes and suggesting new ones if necessary. Finally, we synthesised the codes into three overarching themes, ensuring they align with the research question RQ3.

3.3.2 Assessments. As a part of the assessment-based evaluation, we analysed the frequency of attempts on optional exercises as a proxy for voluntary learning effort. Since these exercises were ungraded, attempt frequency was interpreted as an indicator of students' initiative in reinforcing their privacy knowledge, an essential skill for professional growth.

Then, we evaluated the quality of the MST and final exam questions using the difficulty index (the percentage of students who correctly answered the question) and the discrimination index (performance differentiation between high and low 27% scorers) [25]. The difficulty index (DIF) varies from 0 to 100%, where questions are categorised as easy (>70%), acceptable (30-70%), and difficult (<30%) [26]. The discrimination index (DIS) ranges from -1.00 to +1.00, where negative values indicate non-discriminating questions, while positive values suggest poor (<0.15), marginal (0.15–0.24), good (0.25–0.34), or excellent ($\geq$ 0.35) discrimination [26].

Unlike the optional exercises, the MST, and the final exam, we could not evaluate the assignment. Most groups did not clearly explain individual contributions; as such, we could not confirm whether a considerable number of students had equal exposure to all questions. Consequently, we excluded the assignment results from the evaluation process.

4 Results

This section presents the results of the surveys and assessments. Of the 221 students enrolled, 86 and 139 responded to the pre- and post-surveys, respectively. 35 responses from the post-survey were omitted according to the exclusion criteria EC1 (2), EC2 (14), EC3 (2), EC4 (15), and EC5 (2) mentioned in Section 3.2.

4.1 Learning Progress

During the post-survey, students had to self-report whether they had learned the privacy content before the course. Figure 1 depicts the responses to this question. Overall, students were more familiar with basic concepts related to privacy, such as the meaning of privacy (n = 67), the social aspect of privacy (n = 66), the importance of privacy protection for individuals (n = 74) and organisations (n = 57), and data privacy (n = 62). Students relied primarily on informal learning resources such as social media rather than formal academic materials to learn these topics. Striking unfamiliarity was reported regarding advanced topics such as privacy threats (n = 12), different types of PETs (n = 4) and programming PETs (n = 2), which are crucial for privacy-aware software development.

The responses to the Likert items LP1 to LP6 are visually presented in Appendix C, Figure 5. In the post-survey, most students positively responded to these items compared to the pre-survey. As shown in Table 2, this observed difference was statistically significant ($p < .001$) across all six items, rejecting their null hypotheses of equal distributions. All Likert items except LO2 had a large effect size ($\geq$ 0.14). The Likert item LO2 had a medium effect size (0.06).

The responses to the open-ended question LO7 are depicted as word clouds in Figure 2. These responses suggest that the students struggled to distinguish between privacy and security (e.g., access control, authentication) protection mechanisms. This observation resonates with the confusion between privacy and security among the CS professionals interviewed during privacy research [27, 53]. In the post-survey responses (Figure 2.b), most students could recall the three PETs they programmatically implemented out of the six PETs they learnt. Security measures are still confused with privacy measures in the post-survey responses. However, more diverse privacy protection mechanisms were mentioned. Thus, it is clear that the privacy lessons have instilled a better awareness of privacy protection mechanisms. Overall, the responses from LO1 to LO7 indicate an improvement in the students' knowledge. This finding is also validated by the responses of LO8 (median = 6), which asked the students if the lessons improved their understanding of privacy. LO8 responses are visually presented in Appendix C, Figure 6.

4.2 Perceptions on PETs Programming

This section presents how different learner groups perceived the PETs programming lessons. We extracted the post-survey responses of the Likert items PP1 to PP8 of 81 students who completed at least one programming lesson. The responses are visually presented in Appendix C, Figure 7. As depicted in Table 3, we could group the eight Likert items into three factors following high factor loadings (>0.4). We named these factors Ease of Completing Tasks, Programming Motivation and Applied Learning based on the items grouped under them. According to Cronbach's alpha (0.827, 0.807, 0.785), each factor represented a high internal consistency, indicating that the Likert items under each factor reliably measure the same underlying construct [67]. Next, we inspected how different learners are grouped under these three factors using hierarchical clustering. We stopped the clustering at stage 76 as we observed the first sudden spike in coefficients between stages 76 and 77 [77], resulting in five clusters. The scree plot depicting this scenario and the dendrogram are included in Appendix D, Figure 10.

Following are the descriptions of the five learner groups (clusters) we identified. These groups are summarised in Table 4.

Capable but Unengaged. Compared to the other groups, these learners found it easier to complete the programming tasks (median = 6). While they perceived that programming helps to map privacy theories into practice (median = 5.33), they showed less enthusiasm to engage in programming activities (median = 3.33).

Moderately Engaged Learners. This learner group is the largest among the five (n = 34). Students in this group reported neither difficulty nor ease in completing tasks (median = 4.00). They showed a moderate motivation towards programming PETs (median = 4.33). In contrast, the students demonstrated a notable positive attitude towards applied learning (median = 5.67).

Hands-On Learners. This group easily completed the tasks (median = 5.00) and was highly motivated to participate in programming PETs (median = 6.67). They perceived programming as extremely valuable in understanding theory, with the highest response to applied learning (median = 7.00).

Perseverant Learners. Despite reporting the lowest ease in task completion (median = 2), students in this group maintained a

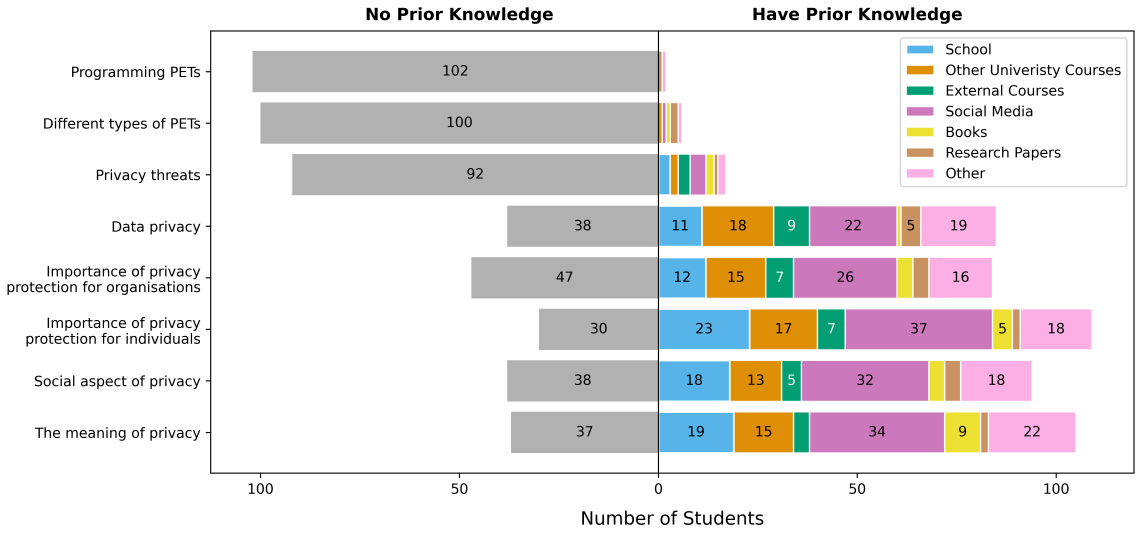

Figure 1: Students' familiarity with different privacy topics before taking the course. The figure also reports the resources students used to gain prior knowledge.

Table 2: Mann–Whitney U test results indicating the learning progress of the students.

	Likert Item	Pre/Post	Mean Rank	U	Z Score	p	η^2 [95% CI]
LP1	What is "Privacy"?	Pre	75.61	2761.50	-5.1614	<.001	0.14 [0.06 - 0.26]
		Post	111.95				
LP2	Rate your awareness regarding privacy threats in the offline world	Pre	82.09	3319.00	-3.4946	<.001	0.06 [0.02 - 0.15]
		Post	106.59				
LP3	Rate your awareness regarding privacy threats in the online world	Pre	73.10	2545.50	-5.7743	<.001	0.18 [0.09 - 0.27]
		Post	114.02				
LP4	Rate your ability to identify privacy vulnerabilities and privacy threats	Pre	70.40	2313.50	-6.1645	<.001	0.20 [0.12 - 0.31]
		Post	116.25				
LP5	Rate your confidence in trying to mitigate the identified privacy threats	Pre	71.83	2436.50	-5.8237	<.001	0.18 [0.09 - 0.28]
		Post	115.07				
LP6	Rate your skill level required to mitigate the identified privacy threats	Pre	71.65	2421.00	-5.7978	<.001	0.18 [0.09 - 0.28]
		Post	115.22				

U = Mann-Whitney U rank, p = 2-tailed significance, η^2 [95% CI] = eta-squared effect size with 95% confidence interval.

relatively high motivation to program PETs (median = 5). They favourably viewed the applied learning experience of the programming tasks (median = 5.33).

Struggling Learners. This small group (n = 7) faced high difficulty in completing tasks (median = 2) and had very low motivation for programming (median = 2). They also reported low perceived value in the active learning experience (median = 3.67).

Figure 3 visually depicts the programming perceptions of each learner group along the three latent factors. We used notched boxplots so that it is possible to visualise the confidence interval (notch) around the medians [72]. Non-overlapping notches in two boxplots indicate a significant difference between their medians. For example, both perseverant and struggling learner groups have overlapping perceptions about the ease of completing the programming tasks. However, their perceptions significantly differ regarding the motivation to program PETs and applied learning.

4.3 Role in Privacy Protection

In analysing how students perceive their involvement in privacy protection, we found that over 75% responded positively to the Likert items PR1 (median = 6), PR2 (median = 6) and PR3 (median = 5.5). The responses of these Likert items are presented in Appendix C, Figure 8. The Cronbach's alpha of the three items is 0.834, indicating a high internal reliability [67]. Next, we analysed how these items correlated with the improved understanding of privacy concepts (LP8). As illustrated in Figure 4, LP8 indicated a strong positive relationship with the perceived relevance (r = 0.595, p <

(a) Pre-survey (b) Post-survey

Figure 2: Responses to the open ended question LP7: "Which mechanisms can be used to protect data privacy during software development?"

Table 3: Principle axis factoring performed on the Likert items that measure the perceptions towards PETs programming. Three factors were identified: Ease of Completing Tasks, Motivation to Program, and Applied Learning.

	Likert Item	Ease of Completing Tasks	Motivation to Program	Applied Learning
PP1	They helped me to understand more about the theory	.102	.256	.814
PP2	It was interesting to convert theory into practice		.202	.770
PP3	They showed me how privacy and software development relates	.131	.269	.644
PP4	They motivated me to program other PETs		.954	.151
PP5	They motivated me to try PETs in my software projects		.679	.375
PP6	I wish I had more exercises		.569	.233
PP7	The exercises were easy	.883		
PP8	Manageable Workload	.736		.223

The Likert items grouped (orange color) under each factor had higher factor loadings (>0.4)

.001) and importance (r = 0.56, p < .001) of the privacy lessons and showed a moderately positive correlation with the perceived responsibility (r = 0.486, p < .001).

Then, we analysed the responses to the open-ended question PR4 to observe how students relate the learnt privacy content to their professional lives. We narrowed the comments to three overarching themes: the *early stage of responsibility, privacy-aware software development, and narrowed perspective on positioning privacy knowledge*. Most responses indicated that the students knew the importance of protecting privacy. However, their level of intended engagement in privacy protection varied. Table 9 in Appendix F includes the codes, themes and example excerpts.

Table 4: Five learner groups based on their perceptions towards PETs programming lessons.

Learner Group (Cluster)	Cluster Size	Median		
		Ease of Completing Tasks	Motivation to Program	Applied Learning
Capable but Unengaged	17	6.00	3.33	5.33
Moderately Engaged	34	4.00	4.33	5.67
Hands-On	7	5.00	6.67	7.00
Perseverant	16	2.50	5.00	5.33
Struggling	7	2.00	2.00	3.67

Early Stage of Responsibility. Some students (n = 14) planned to be more cautious about privacy concerns. However, despite this vigilance, their comments did not display intentions to take proactive steps towards privacy protection.

> *"Now I have a more baseline understanding of privacy importance, and in particular ways its can be threatened"*

> *"The content will make me aware of these issues, and include it in my thinking process"*

> *"I plan to work with technology in the future so privacy will be very important and information/data will need to be protected"*

A few students explicitly showed an interest in protecting their privacy rather than the privacy of others.

> *"be aware of sharing my personal details through everything I interact with"*

> *"When I search for some information on the internet, I will be more careful to share my information"*

Privacy-Aware Software Development. This overarching theme reflects the students' intentions to apply the learnt knowledge in privacy-preserving software development. We further divided this theme into three sub-themes: *privacy-first approach, contribution through programming* and *ownership in privacy protection.*

Privacy-First Approach. Students intended to consider privacy a core element in the software design process rather than an afterthought. Some specifically appreciated the privacy-first proactivity due to the newfound awareness of privacy issues.

> *"When it comes to projects, I think it's a good idea to adopt a 'privacy first' mentality. So when developing programs, privacy will be one of the main considerations to cater to"*

> *"It allowed me to notice more privacy issues in the world around me and I think it will be helpful in the future so that I can incorporate privacy principles into my designs from the beginning stages"*

Two students acknowledged their disagreement with privacy as an afterthought mentality.

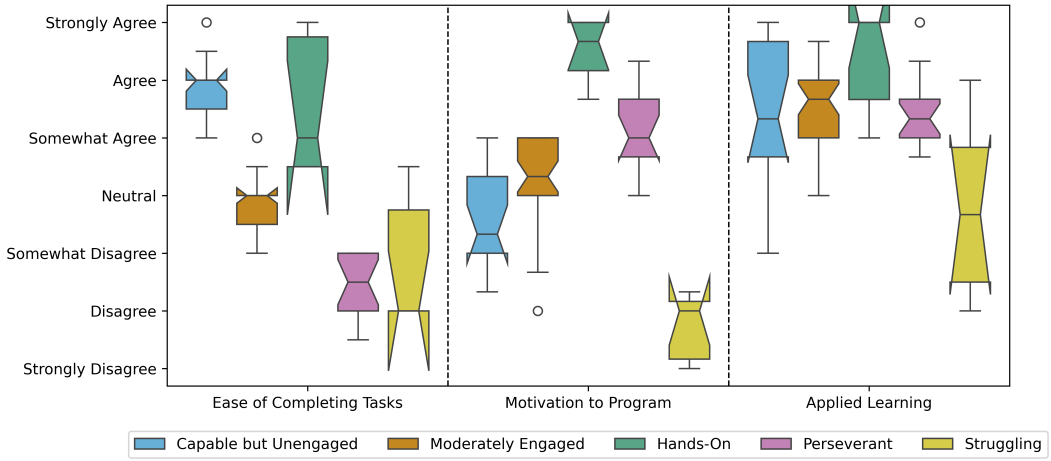

Figure 3: Visual representation of the five learner groups according to their perceptions towards programming PETs. The groups had varying perceptions regarding the ease of completing tasks, motivation to program PETs and applied learning.

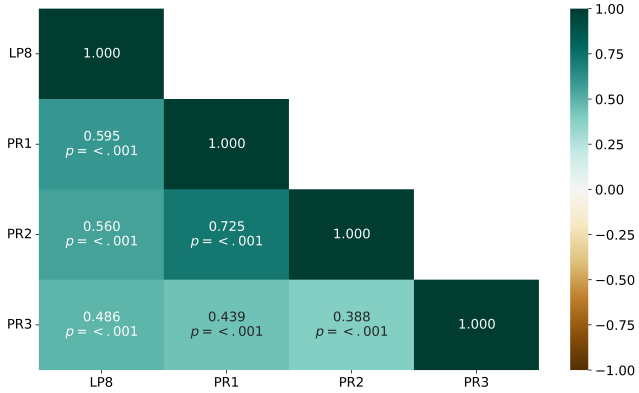

Figure 4: Correlation between PR1 (perceived lesson relevance), PR2 (perceived lesson importance), PR3 (perceived responsibility) and LP8 (privacy understanding).

> *"I'm genuinely surprised that privacy isn't a permanent topic in COMPSCI 316. I think this will definitely help me prioritise privacy in future projects as it is already often forgotten and only implemented when there is a privacy breach"*
>
> *"Being more conscious of building privacy protection measures into software projects rather than as an addition"*

Contribution Through Programming. Several students (n = 8) focused on the programming skills gained through the lessons. They reflected on how the coursework helped reinforce the idea that privacy can be incorporated into programmers' workflow.

> *"assignment 2 made me more familiar with the PETs we covered and the experience implementing them will be good to draw from"*
>
> *"Before week 6 & 7 I had trouble understanding how our knowledge in cyber security could be put into practice in the real world, and getting a glimpse of it through coding different PETs was very helpful in that understanding"*

Furthermore, two students latently conveyed that privacy programming is a crucial skill for professionals, attaching some value to that skill.

> *"I have tried to code for privacy which did make me think about how do real professionals achieve that. So that can be useful for my future career"*
>
> *"An understanding of how privacy can be implemented in an application is helpful, because it should be a key requirement for any developer"*

Ownership in Privacy Protection. A few students (n = 4) expressed their intention to actively contribute towards privacy protection. These comments reflect a strong sense of ownership over privacy protection during software development. Some students planned to apply the learnt knowledge in personal projects, displaying a clear commitment and enthusiasm to taking action.

> *"I will be implementing a PET (K-anonymity) into my own personal project"*
>
> *"Basic understanding of privacy so I can search for more to learn more to implement in work or even further studies :)"*

Additionally, one student reflected on existing privacy issues in a past project and was confident about improving it.

> *"It tells me the importance of protecting privacy. I know why my previous website work is unsafe and knows how to improve it"*

One student shifted from a compliance-driven privacy protection approach to a more internally committed one. Initially, the student followed privacy policies at work out of obligation. However, after engaging with the course content, the student began to grasp the underlying importance of data privacy protection and the rationale

behind the policies and training. The student felt more empowered to take ownership of privacy protection.

> *"As someone who currently works in one of the largest tech companies, we have a lot of policies regarding data privacy protection. I was following them out of compliance, knowing that violations lead to paperwork. The lecture content regarding data privacy helped me understand why we need to protect data privacy in the first place. It has also given me insight into why I have to complete our mandatory training, and helped me understand why there is a need to put these training modules into practice"*

Narrowed Perspective in Positioning Privacy Knowledge. Several students (n = 12) struggled to articulate how to apply the learnt privacy knowledge in their future careers. Some explicitly expressed their uncertainty through comments such as "unsure", "no clue", "I'm not sure if it will", while some questioned the relevance of privacy knowledge to their professional paths.

> *"not much, not that relevant for what I want to do"*
>
> *"I don't think it will be particularly useful for me"*

Despite gaining a better understanding of privacy, one student was reluctant to use the gained knowledge, citing the difficulty in applying it in practice.

> *"Not too much, it improved my understanding but putting it into practice looks too much of a hassle"*

Fragmented understanding of the gained knowledge was also evident. For example, a student expressed a positive view of PETs. However, the response overgeneralised the usage of PETs, suggesting that they could be applied to any form of data handling, which is not necessarily true. The response also reflected a tendency to frame privacy protection primarily in terms of programming PETs.

> *"It was really good to learn about the few common pets and I believe that it'll be useful in any form of data handling. It's quite rare for us to think of what will happen in the future and if we're going to be implementing any of the PETs learnt in class although it's good to have understanding"*

Additionally, several students failed to understand how the learnt knowledge fits diverse professional domains. For example, one associated privacy knowledge primarily with cybersecurity-related professions, and another assumed that privacy concepts were essential for every Information Technology based industry, which may not hold, particularly when personal data is not used.

> *"If I end up working in the field of cybersecurity, it will be very helpful, including that I will be able to communicate with people more calmly in future employment and will not offend them"*
>
> *"Basically, any IT industry will require privacy knowledge."*

4.4 Assessments

The student engagement with the optional exercises is presented in Table 5. The average completion rate of the optional exercises was very low (22.74%). Most students were interested in completing the week 6 practice quiz, which allowed them to evaluate their privacy knowledge before the MST. In addition, even if we expected the students to initiate discussions about the optional exercises, especially for those without answers, we did not observe such interactions on the course's online discussion platform.

Table 5: Student engagement with the optional exercises.

Optional Exercise	Count (n = 221)
Whispers and Echos	65
A Day in Your Life	52
Be Mindful	43
Track Privacy Threats	5
More Than Privacy Protection	11
Which PET(s) Do You Need?	6
Week 6 Practice Quiz	152
Differential Privacy	68

Table 6: The difficulty and discrimination indices of the mid semester and final exam questions.

	Question	DIF	DIS
MST	What type of threat is this (scenario) according to the LINDDUN threat model?	58	0.35
	Select the suitable PETs for the given software scenario	38	0.40
Final Exam	What is incorrect regarding the LINDDUN threat modelling?	62	0.31
	Which of the following is not a key factor to consider when selecting a PET for software applications?	74	0.56
	What is/are the correct conclusion/s that can be taken by analysing the given k-anonymity code?	39	0.32
	Which epsilon has the highest probability of generating the noisy result 1000.17 for the given query if the real value is 1000?	42	0.48

MST = Mid Semester Test, DIF = Difficulty index, DIS = Discrimination index

Next, we evaluated the quality of the MST and final exam questions using item analysis. The results of the difficulty and discrimination indices are displayed in Table 6. Out of all the questions, the second question of the exam was easy (DIF = 74%) for the students. However, it showed excellent discrimination (DIS = 0.56), i.e., higher-performing students often tend to get that question correct than lower-performing students. Students find it difficult to select suitable PETs for a given scenario (DIF = 38%) and answer the programming question (DIF = 39%). However, both questions showed acceptable discrimination indices (0.4 and 0.32).

Despite our expectations, the assignment submissions did not demonstrate significant student collaboration. Only six groups specifically explained their collaborative efforts. 14 groups divided the five questions among the members, while 24 groups failed to specify the individual contributions. Given the limited number of

students who engaged in collaborative work, we could not confidently determine whether all students had equal exposure to all the questions. In the final open-ended question of the post-survey, some students explained how they divided the questions among the group and completed them individually.

> *" I felt as if Assignment 2 being a group assignment hindered my ability to learn, the way the assignment was laid out with 5 questions, and 5 members led to the easy distribution of 1 question per member"*

> *"The questions didn't necessarily feel connected to each other to make it a group assignment, once the group had decided on what question each member was doing, there was no actual groupwork but just 5 students working on their own questions individually and then compiling them for final submission."*

Additionally, one student shared how personal circumstances made it difficult to engage in group work.

> *"As I'm working full time and have 2 kids it's really hard for me to make time in normal student hours to study or collaborate with other students so I didn't want to jeopardize other students grade just because I couldn't make time."*

5 Discussion

In this section, we discuss our findings with reference to the research questions, examine how the curriculum shaped students' privacy mindset, suggest improvements to the proposed curricula and report the limitations of the study.

5.1 How Has the Proposed Curriculum Improved the Privacy Knowledge of Computer Science Undergraduates?

Before taking the privacy lessons, the students noticeably lacked the technical knowledge needed to implement privacy protection in software. They were more familiar with the topics: the meaning of privacy, societal aspects of privacy and the implications of privacy violations. However, most (94.23%) lacked knowledge of privacy threats, PETs and PETs programming. In addition, most students were unaware of different privacy protection mechanisms (including but not limited to PETs) based on the LP7 responses.

According to the pre-survey, the students have learned or heard about privacy mainly from informal resources such as social media rather than formal educational institutions like schools and/or universities. This observation raises questions about the validity and depth of their privacy knowledge before the course. For instance, while the pre-survey responses of LP2 and LP3 indicated that most students were aware of privacy threats, many reported having no prior knowledge of such threats when asked again in the post-survey (Figure 1). A plausible explanation is that the students realised their initial perception of privacy threats was incomplete or inaccurate after learning the LINDDUN model.

Mann-Whitney U test results suggest that the students showed higher self-assessment scores for the learning outcomes after the lessons. This change was noticeably higher in identifying privacy threats and perceived confidence and skill to mitigate privacy threats. This observation is due to learning about privacy threats, PETs and PETs programming, which were noticeable knowledge gaps before the course. However, there is room for improvement in raising students' awareness regarding offline privacy threats, as it had a medium effect size. This awareness is crucial to eliminate the misconception - going offline would protect against privacy threats. For instance, swiping a bank card at a supermarket, a seemingly offline activity, still creates a digital footprint and exposes card owners to online threats [17]. Even though the lectures explained such activities, course assessments heavily focused on identifying privacy threats generated through online activities or software functionalities, thus likely leading to the said knowledge gap.

In addition, while less pronounced than in the pre-survey responses, the knowledge gap of differentiating privacy from security still existed (LP7 responses). We do not attribute this confusion to integrating our curriculum within a cybersecurity course due to several reasons: the curriculum provided explicit definitions of privacy and security, a visual distinction (Venn diagram) in lesson 2 - data privacy protection, and targeted feedback via the assignment (Question 3.3). However, the indirect role of security in privacy protection and low retention of the students might have caused the confusion. Thus, curriculum improvements that frequently reinforce the security and privacy distinction and clearly phrased survey questions that avoid ambiguity are required.

Key Takeaways

Practical components (e.g., threat modelling, PETs - theory and programming) should be more emphasised during curricula design as those knowledge are essential in professional settings and are rarely acquired through informal resources. Other CS courses addressing data protection (e.g., database) can also benefit from our lessons, as ACM/IEEE standards endorse such inclusions [36]. Research should explore tools, teaching methods, and curricula that lower the barriers to teaching privacy-focused programming. Lessons should also clarify nuanced topics like the distinction between privacy and security, where confusion can often arise.

5.2 How Do Computer Science Undergraduates Perceive Privacy Programming Lessons?

The clustering analysis revealed five distinct learner groups based on their perceptions towards PETs programming: *capable but unengaged*, *moderately engaged*, *hands-on*, *perseverant* and *struggling*. All groups except the struggling group (8.6%) recognised the value of the programming exercises in bridging the gap between theory and practice. Thus, the majority agrees that practical knowledge is essential to truly grasp privacy concepts, i.e., 'do it to know it'. In this context, the proposed experiential exercises contribute to easing the pressure of privacy regulations, which often instruct what to do but not how to do when it comes to privacy [34, 58, 61].

Unlike the perception towards applied learning, task completion difficulty and motivation to program PETs indicated noticeable differences among the groups. The ease of completing tasks did not consistently predict the motivation to program PETs. The perseverant group struggled to complete the tasks yet remained highly

motivated to engage with PETs programming. If task difficulty and workload were adjusted to be more comfortable, these students could transition into the hands-on learner group, where the motivation to program and ease of task completion are both high.

Conversely, capable but unengaged students demonstrated a low motivation to program PETs despite finding the tasks easy to complete. This observation is possible if those students' interests and career aspirations are not aligned with programming activities. Reducing the difficulty and workload of the programming tasks alone would not be enough to increase their programming motivation. Moderately engaged students, the largest group (41.9%), did not find the tasks easy or difficult, yet they still saw value in engaging with them. If task difficulty and workload were fine-tuned, some students in this group could shift toward higher motivation levels, similar to those in the hands-on learner group.

Key Takeaways

The difficulty level and design of privacy programming exercises should be tailored to meet various learner needs. This flexibility would support students with diverse programming abilities and varying levels of interest in coding. Other programming courses can also use our learner group analysis method to design more inclusive and engaging exercises. Privacy educators can readily adapt our shared programming exercises to suit diverse instructional contexts.

5.3 How Can the Proposed Curriculum Help Computer Science Undergraduates Realise Their Role in Privacy Protection?

According to the post-survey responses, most students believed that privacy lessons are relevant to CS students and that these lessons should be continued for undergraduates. This finding reflects that the students recognise the need for professional competence in resolving privacy matters. The lessons also improved their sense of responsibility in privacy protection. However, the thematic analysis showed that this perceived responsibility had different levels.

Some students were vigilant yet did not indicate the intention to use proactive measures. We believe this is a transitional stage, where students may require further guidance to realise a stronger perception of responsibility. Conversely, some students intended to take proactive steps at work in the future (e.g., adopting privacy-first design and PETs programming), and some exhibited a strong internal commitment to protect privacy (e.g., planning to integrate privacy concepts into personal projects). These varying levels of perceived responsibility can be due to personal traits, such as hope of success (expectation to achieve success) and locus of control (the extent to which individuals believe they have control over the outcomes of events) [65]. For instance, students who intend to follow a privacy-first development approach have hope of success regarding proactive measures, while students who plan to apply the newfound knowledge in personal projects possess both hope of success and a strong locus of control.

Further, some students struggled to apply the learned privacy concepts to their future careers. It is possible that they assumed the lessons focused on future software developers due to the programming lessons. Alternatively, their career aspirations may fall outside the CS field. This sense of exclusion can also affect their perceived responsibility as they no longer see their professional contribution towards privacy protection. Therefore, adapting the lessons to show how privacy knowledge can be applied across different career paths is important. For example, the lessons could show how roles other than developers can apply their understanding of PETs to choose privacy-conscious tools.

Key Takeaways

Privacy curricula should emphasise the multidisciplinary nature of privacy protection. This will foster a sense of inclusion in students, as they realise they can play a meaningful role in privacy concerns regardless of their career aspirations. Other CS courses can also use such reflection strategies to help students connect their learning to real-world contexts.

5.4 Impact on the Privacy Mindset

Mindset refers to the individuals' belief system, which shapes how they interpret information, make decisions, and approach future behaviour [13]. Mindsets are changeable, and learning experience is one way to support this change [13]. As such, the learning experience from our curriculum showed evidence of changing the privacy mindset of the students.

As discussed in Section 5.1, students demonstrated knowledge gains across multiple topic areas. This new knowledge reshaped their beliefs about what privacy is and why it matters. Responses to the open-ended question PR4 (Appendix F, Table 9) also reflect this shift. For instance, one student remarked, *"Now I have a more baseline understanding of privacy importance, and in particular ways it can be threatened"*, expressing a clear privacy mindset shift.

Hands-on exercises played a notable role in shaping students' privacy mindsets. Section 4.3 discussed how some students viewed privacy as an actionable and achievable ethic in their future professional work. They believed privacy should guide their decisions during system design and development. This attitude reflects a shift toward taking ownership of privacy protection and adopting a privacy-first approach to software development. Thus, the curriculum was able to direct students from abstract awareness to a practical, value-driven stance toward privacy protection.

Overall, the curriculum encouraged most students to reevaluate their beliefs on privacy, their role in upholding it, and their capability to integrate it into their future work. Though a few students expressed frustration with the technical implementation and were uncertain about applying the learnt knowledge, we believe curriculum improvements (Section 5.5) may help reduce such attitudes and influence a more confident, actionable privacy mindset.

5.5 Improvements to the Curriculum

We recommend the following improvements to the proposed curriculum based on the evaluation results.

New Content. We suggest three new content areas to deepen the understanding of privacy: adversarial perspective, privacy requirement generation, and knowledge application across different professions.

Adversarial thinking helps students anticipate potential privacy threats and reinforces the need for privacy protection [46]. Example exercises include analysing why specific systems are targeted for privacy breaches, explaining attacks such as website fingerprinting, and studying real-world privacy breaches. Second, privacy requirements generation will assist students in identifying privacy needs early in the development process [8, 27]. Lessons can teach how to combine the results of threat modelling, customer requirements, and requirements drawn from privacy regulations to compose a comprehensive list of requirements. Finally, the course could introduce a section to explain how the learnt knowledge can be applied across different professions. Even in software development, students need to understand how various roles, such as project managers, business analysts, quality analysts, and even higher management, can contribute to protecting privacy. The students who decide to pursue a profession outside the software industry can leverage the learnt knowledge to make informed choices about privacy-aware tools, safeguard their data, and protect the privacy of employees or customers if applicable. By highlighting these diverse applications, the course could help CS undergraduates view privacy as a fundamental consideration regardless of their chosen careers.

Privacy vs Security. The lessons and exercises should reinforce this distinction more clearly and frequently. For example, the curriculum can include real-world scenarios that highlight the difference, compare security and privacy requirements, and compare security and privacy protection techniques.

Programming Lessons. The proposed curriculum had few practice exercises for PETs programming. Providing more hands-on exercises to practice could reduce the perceived difficulty. Following game design theories, these exercises could gradually progress through different levels of challenge, i.e., easy to difficult, to increase engagement despite learners' skills [7]. This strategy will help students gradually improve their programming skills. Further, the course could include reflective questions during the programming tasks to improve students' reasoning abilities, understand the implications of privacy programming and map theoretical knowledge with practical skills. For example, one such question could be: "What privacy threat, according to the LINDDUN method, might arise if k-anonymity is not used in this scenario?".

Assessments. According to the results, students were reluctant to complete the optional exercises. This observation can likely be attributed to two factors. First, students may have felt disengaged due to the absence of immediate rewards. This assumption is supported by the high completion rate of the 'Week 6 practice quiz', which targeted the mid-semester exam. Second, the students might have felt the essay questions required more effort and time. To address these issues, the course could award bonus marks for the optional exercises and maintain a leaderboard for peer-driven motivation [56]. Further, including diverse questions, such as fill-in-the-blanks and capture-the-flag challenges, could reduce monotony. [70].

In addition, the group size could be reduced to two or three persons to make the group assignment more manageable. This will also ensure that students engage meaningfully with all the questions. Further, more exploratory questions instead of simple recall or information lookup questions could be included to improve the individual gain during the group work [75].

5.6 Limitations

This study is subject to several limitations. First, the study was conducted at a single institution and in one iteration, constraining the findings' generalisability. However, curriculum contents are generalisable as they were developed in response to industry demands, aligned with curricular standards [16, 36], and designed using well-grounded privacy research (Appendix A) and Bloom's taxonomy. Further, our transparent and in-depth evaluation supports methodological reproducibility across different contexts. Second, self-reported survey responses are subject to biases and misinterpretations. Nevertheless, we filtered out invalid responses using the exclusion criteria to the best of our ability. In the future, we aim to improve the survey design by avoiding double-barrelled and ambiguous questions, eliminating leading statements, and adding attention-check questions [44]. For example, the security and privacy confusion observed in LP7 could have been reduced if the question was rephrased as "Which mechanisms can be used during software development to **primarily** protect data privacy?". Third, the thematic analysis results are subjective to researcher bias. However, incorporating all authors in the coding and theme development reduces it. Finally, the results are influenced by the quality of the lessons and teaching methods. We hope that the suggested curriculum improvements will lead to a higher curriculum quality.

6 Conclusion

This paper presented the design and evaluation of an undergraduate CS privacy curriculum. It was designed to address the challenges computing professionals face when embedding privacy into software. The curriculum lessons follow a logical progression: teaching the importance of privacy protection, then systematically identifying privacy threats using the LINDDUN model and introducing six PETs, with hands-on programming experience in three of them (pseudonymisation, k-anonymity, and differential privacy). Overall, the students demonstrated an improvement across the defined learning outcomes. While most valued the applied learning experience of programming, they had noticeable differences in motivation to program and perceived task completion difficulty. We identified five distinct learner personalities based on these perceptions towards the programming lessons. In addition, most students acknowledged the relevance and importance of the proposed content for CS undergraduates. However, they showed varying levels of responsibility towards privacy protection. Some recognised the need for privacy but showed little intention to take proactive measures, some demonstrated a strong commitment and ownership of privacy protection, and a few students struggled to relate the learned content to their future careers. Drawing from these findings, we discussed the privacy mindset changes of the students, followed by suggestions to improve the curriculum.

Acknowledgments

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. We thank the anonymous reviewers for their constructive feedback, which helped improve the paper. We also appreciate Andrew Luxton-Reilly and the Computing Education Research Group at the University of Auckland for their suggestions during the early stages of this work. We further thank Robert Biddle, Jun Seo, and Lasini Liyanage for their valuable feedback on the paper.

References

[1] Nitin Agrawal, Reuben Binns, Max Van Kleek, Kim Laine, and Nigel Shadbolt. 2021. Exploring Design and Governance Challenges in the Development of Privacy-Preserving Computation. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 68, 13 pages. https://doi.org/10.1145/3411764.3445677

[2] Sami Alkhatib, Jenny Waycott, George Buchanan, Marthie Grobler, and Shuo Wang. 2021. Privacy by Design in Aged Care Monitoring Devices? Well, Not Quite Yet!. In *Proceedings of the 32nd Australian Conference on Human-Computer Interaction (OzCHI '20)*. Association for Computing Machinery, New York, NY, USA, 492–505. https://doi.org/10.1145/3441000.3441049

[3] Noura Alomar and Serge Egelman. 2022. Developers say the darnedest things: Privacy compliance processes followed by developers of child-directed apps. *Proc. Priv. Enhancing Technol.* 2022, 4 (Oct. 2022), 250–273.

[4] Heba Aly, Yizhou Liu, Reza Ghaiumy Anaraky, Sushmita Khan, Moses Namara, Kaileigh Angela Byrne, and Bart Knijnenburg. 2024. Tailoring Digital Privacy Education Interventions for Older Adults: A Comparative Study on Modality Preferences and Effectiveness. *Proceedings on Privacy Enhancing Technologies* 2024 (Jan. 2024), 635–656. https://doi.org/10.56553/popets-2024-0036

[5] Renana Arizon-Peretz, Irit Hadar, Gil Luria, and Sofia Sherman. 2021. Understanding developers' privacy and security mindsets via climate theory. *Empir Software Eng* 26, 6 (Nov 2021), 123. https://doi.org/10.1007/s10664-021-09995-z

[6] Sumit Asthana, Jane Im, Zhe Chen, and Nikola Banovic. 2024. "I know even if you don't tell me": Understanding Users' Privacy Preferences Regarding AI-based Inferences of Sensitive Information for Personalization. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) *(CHI '24)*. Association for Computing Machinery, New York, NY, USA, Article 782, 21 pages. https://doi.org/10.1145/3613904.3642180

[7] Barbaros Bostan and Sertaç Öğüt. 2009. Game challenges and difficulty levels: lessons learned From RPGs. In *International simulation and gaming association conference*.

[8] Maisha Boteju, Thilina Ranbaduge, Dinusha Vatsalan, and Nalin Arachchilage. 2024. SoK: Demystifying Privacy Enhancing Technologies Through the Lens of Software Developers. (April 2024). https://doi.org/10.48550/arXiv.2401.00879

[9] Virginia Braun and Victoria Clarke. 2021. *Thematic analysis*. SAGE Publications, London, England.

[10] Nicole M. Cain, Callie Jowers, Mark Blanchard, Sharon Nelson, and Steven K. Huprich. 2021. Examining the Interpersonal Profiles and Nomological Network Associated with Narcissistic Grandiosity and Narcissistic Vulnerability. *Psychopathology* 54, 1 (2021), 26–38. https://doi.org/10.1159/000510475

[11] Ann Cavoukian. 2009. Privacy by Design.The Answer to Overcoming Negative Externalities Arising from Poor Management of Personal Data. In *Trust Economics Workshop London, England, June*, Vol. 23. 2009.

[12] Roger Clarke. 1997. Introduction to Dataveillance and Information Privacy, and Definitions of Terms. https://www.rogerclarke.com/DV/Intro.html

[13] Carol S. Dweck. 2006. *Mindset: the new psychology of success* (1st ed ed.). Random House, New York.

[14] Cynthia Dwork, Nitin Kohli, and Deirdre Mulligan. 2019. Differential Privacy in Practice: Expose your Epsilons! *Journal of Privacy and Confidentiality* 9, 2 (Oct. 2019). https://doi.org/10.29012/jpc.689

[15] Cynthia Dwork and Aaron Roth. 2014. The Algorithmic Foundations of Differential Privacy. *Found. Trends Theor. Comput. Sci.* 9, 3–4 (Aug. 2014), 211–407. https://doi.org/10.1561/0400000042

[16] Joint Task Force On Cybersecurity Education. 2018. *Cybersecurity Curricula 2017: Curriculum Guidelines for Post-Secondary Degree Programs in Cybersecurity*. ACM, New York, NY, USA. https://doi.org/10.1145/3184594

[17] Serge Egelman, Julia Bernd, Gerald Friedland, and Dan Garcia. 2016. The Teaching Privacy Curriculum. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education* (Memphis, Tennessee, USA) *(SIGCSE '16)*. Association for Computing Machinery, New York, NY, USA, 591–596. https://doi.org/10.1145/2839509.2844619

[18] Anirudh Ekambaranathan, Jun Zhao, and Max Van Kleek. 2021. "Money Makes the World Go around": Identifying Barriers to Better Privacy in Children's Apps From Developers' Perspectives. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 46, 15 pages. https://doi.org/10.1145/3411764.3445599

[19] European Commission. 2016. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance). https://eur-lex.europa.eu/eli/reg/2016/679/oj

[20] Rachel L. Finn, David Wright, and Michael Friedewald. 2013. *Seven Types of Privacy*. Springer Netherlands, Dordrecht, 3–32. https://doi.org/10.1007/978-94-007-5170-5_1

[21] UN System Chief Executives Board for Coordination. N/A. Principles on Personal Data Protection and Privacy. https://unsceb.org/principles-personal-data-protection-and-privacy-listing

[22] European Union Agency for Cybersecurity. 2021. *Data pseudonymisation: advanced techniques and use cases*. Publications Office. https://data.europa.eu/doi/10.2824/860099

[23] Sarah Abdelwahab Gaballah, Lamya Abdullah, Ephraim Zimmer, Sascha Fahl, Max Mühlhäuser, and Karola Marky. 2025. "It's Not My Data Anymore": Exploring Non-Users' Privacy Perceptions of Medical Data Donation Apps. *Proceedings on Privacy Enhancing Technologies* 2025, 1 (Jan. 2025), 654–670. https://doi.org/10.56553/popets-2025-0035

[24] Viktor Hargitai, Irina Shklovski, and Andrzej Wasowski. 2018. Going Beyond Obscurity: Organizational Approaches to Data Anonymization. *Proc. ACM Hum.-Comput. Interact.* 2, CSCW, Article 66 (Nov. 2018), 22 pages. https://doi.org/10.1145/3274335

[25] SM Hetzel. 1997. Basic concepts in item and test analysis. *EricDatabase, Retrieved July* 25 (1997), 2000.

[26] Mozaffer Rahim Hingorjo and Farhan Jaleel. 2012. Analysis of one-best MCQs: the difficulty index, discrimination index and distractor efficiency. *JPMA-Journal of the Pakistan Medical Association* 62, 2 (2012), 142.

[27] Stefan Albert Horstmann, Samuel Domiks, Marco Gutfleisch, Mindy Tran, Yasemin Acar, Veelasha Moonsamy, and Alena Naiakshina. 2024. "Those things are written by lawyers, and programmers are reading that." Mapping the Communication Gap Between Software Developers and Privacy Experts. *Proceedings on Privacy Enhancing Technologies* 2024, 1 (Jan. 2024), 151–170. https://doi.org/10.56553/popets-2024-0010

[28] ISO. 2019. Security techniques — Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management — Requirements and guidelines (ISO/IEC 27701:2019). https://www.iso.org/standard/71670.html

[29] Leonardo Horn Iwaya, Muhammad Ali Babar, and Awais Rashid. 2023. Privacy Engineering in the Wild: Understanding the Practitioners' Mindset, Organizational Aspects, and Current Practices. *IEEE Transactions on Software Engineering* 49, 9 (2023), 4324–4348. https://doi.org/10.1109/TSE.2023.3290237

[30] Cohen Jacob. 1988. *Statistical Power Analysis for the Behavioral Sciences*. Vol. 2. New York: Academic Press. https://doi.org/10.1016/C2013-0-10517-X

[31] Dali Kaafar, Hassan Asghar, and Raghav Bhaskar. 2021. *From probable to provable privacy*. Retrieved January 27, 2025 from https://www.mq.edu.au/partner/access-business-opportunities/innovation-entrepreneurship-and-it/cyber-security-hub/news/news/from-probable-to-provable-privacy

[32] Miloš Kankaraš and Stefania Capecchi. 2024. Neither agree nor disagree: use and misuse of the neutral response category in Likert-type scales. *METRON* (Sept. 2024). https://doi.org/10.1007/s40300-024-00276-5

[33] Alexandra Klymenko, Stephen Meisenbacher, Iva Lilova, and Florian Matthes. 2024. Investigating the Motivational Factors Influencing Managerial Decisions to Adopt Privacy-Enhancing Technologies. *ECIS 2024 Proceedings* (2024).

[34] Oleksandra Klymenko, Oleksandr Kosenkov, Stephen Meisenbacher, Parisa Elahidoost, Daniel Mendez, and Florian Matthes. 2022. Understanding the Implementation of Technical Measures in the Process of Data Privacy Compliance: A Qualitative Study. In *Proceedings of the 16th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement* (Helsinki, Finland) *(ESEM '22)*. Association for Computing Machinery, New York, NY, USA, 261–271. https://doi.org/10.1145/3544902.3546234

[35] David R. Krathwohl. 2002. A Revision of Bloom's Taxonomy: An Overview. *Theory Into Practice* 41, 4 (2002), 212–218. http://www.jstor.org/stable/1477405

[36] Amruth N. Kumar, Rajendra K. Raj, Sherif G. Aly, Monica D. Anderson, Brett A. Becker, Richard L. Blumenthal, Eric Eaton, Susan L. Epstein, Michael Goldweber, Pankaj Jalote, Douglas Lea, Michael Oudshoorn, Marcelo Pias, Susan Reiser, Christian Servin, Rahul Simha, Titus Winters, and Qiao Xiang. 2024. *Computer Science Curricula 2023*. ACM, New York, NY, USA. https://doi.org/10.1145/3664191

[37] K. LeFevre, D.J. DeWitt, and R. Ramakrishnan. 2006. Mondrian Multidimensional K-Anonymity. In *22nd International Conference on Data Engineering (ICDE'06)*. 25–25. https://doi.org/10.1109/ICDE.2006.101

[38] Na Li, Vikram Chava, and Lin Li. 2017. A labware for educating location privacy protection in location-based services. *J. Comput. Sci. Coll.* 32, 4 (April 2017), 40–48.
[39] Tianshi Li, Elizabeth Louie, Laura Dabbish, and Jason I. Hong. 2021. How Developers Talk About Personal Data and What It Means for User Privacy: A Case Study of a Developer Forum on Reddit. *Proc. ACM Hum.-Comput. Interact.* 4, CSCW3, Article 220 (jan 2021), 28 pages. https://doi.org/10.1145/3432919
[40] Ze Shi Li, Colin Werner, Neil Ernst, and Daniela Damian. 2022. Towards privacy compliance: A design science study in a small organization. *Information and Software Technology* 146 (2022), 106868. https://doi.org/10.1016/j.infsof.2022.106868
[41] Ying Ma, Cherie Sew, Zhanna Sarsenbayeva, Jarrod Knibbe, and Jorge Goncalves. 2024. Understanding Users' Perspectives on Location Privacy Management on iPhones. *Proc. ACM Hum.-Comput. Interact.* 8, MHCI, Article 282 (Sept. 2024), 25 pages. https://doi.org/10.1145/3676529
[42] Miriam Magdolen, Sascha Von Behren, Jan Vallée, Bastian Chlond, and Peter Vortisch. 2024. Response bias in Likert-style psychological items – an example from a large-scale travel survey in China. *Transportation Research Procedia* 76 (2024), 349–360. https://doi.org/10.1016/j.trpro.2023.12.060
[43] Henry B. Mann and Douglas R. Whitney. 1947. On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *Annals of Mathematical Statistics* 18 (1947), 50–60. https://api.semanticscholar.org/CorpusID:14328772
[44] Brian M. McSkimming, Sean Mackay, and Adrienne Decker. 2021. Investigating the usage of Likert-style items within Computer Science Education Research Instruments. In *2021 IEEE Frontiers in Education Conference (FIE)*. 1–8. https://doi.org/10.1109/FIE49875.2021.9637198
[45] J. McTighe and G.P. Wiggins. 1999. *Understanding by Design Handbook*. Association for Supervision and Curriculum Development. https://books.google.co.nz/books?id=bp8DAAAACAAJ
[46] Marina Moore, Maximilian Zinkus, Nathan Lemay, Zachary Peterson, and Bruce DeBruhl. 2018. Introducing privacy to undergraduate computing students. *J. Comput. Sci. Coll.* 33, 4 (April 2018), 157–164.
[47] Gonzalo Munilla Garrido, Xiaoyuan Liu, Floria Matthes, and Dawn Song. 2023. Lessons Learned: Surveying the Practicality of Differential Privacy in the Industry. *Proceedings on Privacy Enhancing Technologies* 2023, 2 (April 2023), 151–170. https://doi.org/10.56553/popets-2023-0045
[48] Nadim Nachar. 2008. The Mann-Whitney U: A Test for Assessing Whether Two Independent Samples Come from the Same Distribution. *Tutorials in Quantitative Methods for Psychology* 4, 1 (March 2008), 13–20. https://doi.org/10.20982/tqmp.04.1.p013
[49] Joseph Near and Chiké Abuah. 2025. *Programming Differential Privacy*. N/A.
[50] Helen Nissenbaum. 2010. Privacy in Context: Technology, Policy, and the Integrity of Social Life. *Bibliovault OAI Repository, the University of Chicago Press* (Jan. 2010).
[51] Marie Caroline Oetzel and Sarah Spiekermann. 2014. A systematic methodology for privacy impact assessments: a design science approach. *European Journal of Information Systems* 23, 2 (Mar 2014), 126–150. https://doi.org/10.1057/ejis.2013.18
[52] Information Commissioner's Office. 2023. Privacy-enhancing technologies (PETs). https://ico.org.uk/media/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/privacy-enhancing-technologies-1-0.pdf
[53] Mariana Peixoto, Dayse Ferreira, Mateus Cavalcanti, Carla Silva, Jéssyka Vilela, João Araújo, and Tony Gorschek. 2023. The Perspective of Brazilian Software Developers on Data Privacy. *J. Syst. Softw.* 195, C (Jan 2023). https://doi.org/10.1016/j.jss.2022.111523
[54] Nithin Prabhu. 2018. K-Anonymity. https://github.com/Nuclearstar/K-Anonymity/blob/master/k-Anonymity.ipynb
[55] Marie-Therese Puth, Markus Neuhäuser, and Graeme D. Ruxton. 2015. Effective use of Spearman's and Kendall's correlation coefficients for association between two measured traits. *Animal Behaviour* 102 (2015), 77–84. https://doi.org/10.1016/j.anbehav.2015.01.010
[56] Elias Ratinho and Cátia Martins. 2023. The role of gamified learning strategies in student's motivation in high school and higher education: A systematic review. *Heliyon* 9, 8 (2023), e19033. https://doi.org/10.1016/j.heliyon.2023.e19033
[57] Brandt Redd, Ying Tang, Hadar Ziv, and Sameer Patil. 2024. Layering Sociotechnical Cybersecurity Concepts Within Project-Based Learning. In *Proceedings of the 2024 ACM Conference on International Computing Education Research - Volume 1 (ICER '24)*. Association for Computing Machinery, New York, NY, USA, 406–418. https://doi.org/10.1145/3632620.3671093 event-place: Melbourne, VIC, Australia.
[58] Conrad Sanderson, David Douglas, Qinghua Lu, Emma Schleiger, Jon Whittle, Justine Lacey, Glenn Newnham, Stefan Hajkowicz, Cathy Robinson, and David Hansen. 2023. AI ethics principles in practice: Perspectives of designers and developers. *IEEE Transactions on Technology and Society* (2023).
[59] Jayshree Sarathy, Sophia Song, Audrey Haque, Tania Schlatter, and Salil Vadhan. 2023. Don't Look at the Data! How Differential Privacy Reconfigures the Practices of Data Science. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems (CHI '23)*. Association for Computing Machinery, New York, NY, USA. https://doi.org/10.1145/3544548.3580791
[60] Security and Privacy Academy. 2023. Title. https://www.youtube.com/watch?v=aMDXM5HlB64
[61] Awanthika Senarath and Nalin Asanka Gamagedara Arachchilage. 2018. Understanding software developers' approach towards implementing data Minimization. *The 4th Workshop on Security Information Workers (WSIW), 14th Symposium on Usability, Privacy, and Security, USENIX* (Aug. 2018).
[62] Ben Rydal Shapiro, Amanda Meng, Cody O'Donnell, Charlotte Lou, Edwin Zhao, Bianca Dankwa, and Andrew Hostetler. 2020. Re-Shape: A Method to Teach Data Ethics for Data Science Education. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. https://doi.org/10.1145/3313831.3376251 event-place: Honolulu, HI, USA.
[63] Katie Shilton, Donal Heidenblad, Adam Porter, Susan Winter, and Mary Kendig. 2020. Role-Playing Computer Ethics: Designing and Evaluating the Privacy by Design (PbD) Simulation. *SCIENCE AND ENGINEERING ETHICS* 26, 6 (Dec. 2020), 2911–2926. https://doi.org/10.1007/s11948-020-00250-0
[64] Varun Shiri, Maggie Xiong, Jinghui Cheng, and Jin L.C. Guo. 2024. Motivating Users to Attend to Privacy: A Theory-Driven Design Study. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference* (Copenhagen, Denmark) *(DIS '24)*. Association for Computing Machinery, New York, NY, USA, 258–275. https://doi.org/10.1145/3643834.3661544
[65] Sarah Spiekermann, Jana Korunovska, and Marc Langheinrich. 2019. Inside the Organization: Why Privacy and Security Engineering Is a Challenge for Engineers. *Proc. IEEE* 107, 3 (2019), 600–615. https://doi.org/10.1109/JPROC.2018.2866769
[66] Latanya Sweeney. 2002. k-anonymity: a model for protecting privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.* 10, 5 (Oct. 2002), 557–570. https://doi.org/10.1142/S0218488502001648
[67] Keith S. Taber. 2018. The Use of Cronbach's Alpha When Developing and Reporting Research Instruments in Science Education. *Research in Science Education* 48, 6 (Dec. 2018), 1273–1296. https://doi.org/10.1007/s11165-016-9602-2
[68] Mohammad Tahaei, Alisa Frik, and Kami Vaniea. 2021. Privacy Champions in Software Teams: Understanding Their Motivations, Strategies, and Challenges. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Yokohama, Japan) *(CHI '21)*. Association for Computing Machinery, New York, NY, USA, Article 693, 15 pages. https://doi.org/10.1145/3411764.3445768
[69] Maciej Tomczak and Ewa Tomczak. 2014. The need to report effect size estimates revisited. An overview of some recommended measures of effect size. *Trends in Sport Sciences* 21 (2014), 19 –25. Issue 1.
[70] Wolfgang Vigl and Svetlana Abramova. 2024. Design and Use of Privacy Capture-the-Flag Challenges in an Introductory Class on Information Privacy and Security. In *Proceedings of the 2024 on Innovation and Technology in Computer Science Education V. 1* (Milan, Italy) *(ITiCSE 2024)*. Association for Computing Machinery, New York, NY, USA, 618–624. https://doi.org/10.1145/3649217.3653572
[71] Joe H. Ward. 1963. Hierarchical Grouping to Optimize an Objective Function. *J. Amer. Statist. Assoc.* 58, 301 (March 1963), 236–244. https://doi.org/10.1080/01621459.1963.10500845
[72] James B. Wells and Ben H. Layne. 1996. The Effects of Display Characteristics on the Bias of Estimates of Whisker Length of Regular and Notched Boxplots. *Journal of Educational and Behavioral Statistics* 21, 3 (1996), 247. https://doi.org/10.2307/1165271
[73] Rick Wicklin. 2023. Weak or strong? How to interpret a Spearman or Kendall correlation. https://blogs.sas.com/content/iml/2023/04/05/interpret-spearman-kendall-corr.html
[74] Frank Wilcoxon. 1945. Individual Comparisons by Ranking Methods. *Biometrics Bulletin* 1, 6 (Dec. 1945), 80. https://doi.org/10.2307/3001968
[75] Kristy J. Wilson, Peggy Brickman, and Cynthia J. Brame. 2018. Group Work. *CBE Life Sciences Education* 17, 1 (2018), fe1. https://doi.org/10.1187/cbe.17-12-0258
[76] Kim Wuyts and Wouter Joosen. 2015. LINDDUN privacy threat modeling: a tutorial. *CW Reports* (2015).
[77] Odilia Yim and Kylee T. Ramdeen. 2015. Hierarchical Cluster Analysis: Comparison of Three Linkage Measures and Application to Psychological Data. *The Quantitative Methods for Psychology* 11, 1 (Feb. 2015), 8–21. https://doi.org/10.20982/tqmp.11.1.p008
[78] Hongyi Zhang, Anas Dakkak, David Issa Mattos, Jan Bosch, and Helena Holmström Olsson. 2021. Towards Federated Learning: A Case Study in the Telecommunication Domain. In *Software Business*, Xiaofeng Wang, Antonio Martini, Anh Nguyen-Duc, and Viktoria Stray (Eds.). Springer International Publishing, Cham, 238–253.

A The Lessons

A.1 Introduction to Privacy

Without a clear understanding of privacy, it is difficult to establish its scope, leading to misinterpretations such as the blurred boundary between security and privacy. Thus, this lesson first presented a

definition of privacy proposed by Roger Clarke [12]. This definition is easy to understand and relatable to the students' lives.

> *"Privacy is the interest that individuals have in sustaining a 'personal space', free from interference by other people and organisations"*

For the definition to be more relatable, we showed the students real-life examples that match the notion of privacy (e.g., blinds in a room, privacy screen on a phone). The lesson also presented seven dimensions of privacy: body, behaviour, personal communications, personal data (data privacy), thoughts and feelings, location and space, and relationships [12, 20]. The intention of presenting these dimensions was to break the misunderstanding that privacy is solely about data. Then, the lesson provided how context matters to privacy [50]. Example questions such as "Can you share a conversation you had with your close friends with the whole class?" were asked to familiarise the students with "contextual privacy".

Next, the lessons helped students recognise the importance of protecting privacy. First, we explained the social aspect of privacy in two ways: 1. privacy is shaped by society and 2. privacy empowers society. Second, we explained why privacy is becoming a hot topic by drawing their attention to the exponential data collection and handling in the evolving digital landscape. Third, we discussed how offline actions, such as tapping a transportation card on a bus or paying in a supermarket, can generate online data. These scenarios showed the students that they cannot escape privacy threats by going offline. Finally, we explained the consequences of privacy breaches on individuals and organisations.

A.2 Data Privacy Protection

Computer professionals engage tightly with data privacy when they develop software that handles personal data. Therefore, in this lesson, we narrowed the scope of the lectures to data privacy so that the content would be more relevant to the CS undergraduates.

> *"Data privacy is the control individuals have to decide when they are disclosing personal data, to whom they are disclosing those data and how much data they are disclosing"* [12, 20]

The lesson explained the meaning of personal data using the definition given in the General Data Protection Regulation [19].

> *"Personal data are data related to an identified or identifiable natural person"*

However, we pointed out the vagueness of the definition (no clear boundaries) and then justified why such vagueness is required due to the context dependency in privacy. Also, we explained how such vagueness can influence organisations to be more cautious when handling personal data.

We then explained the concept of 'digital footprint' and how it is generated through intentionally shared data, hidden data, online activities, data shared by others, and inferred data [17]. Following this, we taught students how to protect their own data in the digital world (e.g., cookie blockers, incognito mode).

Next, we presented a case study, "Facebook Cambridge-Analytica incident", to prove how poor programming practices and human errors of the data collectors can lead to data breaches. Following this explanation, we explained how such instances can be mitigated using Privacy by Design [11] (a set of guidelines that help to consider privacy concerns proactively) and the LINDDUN threat model [76] (a framework that helps to identify different types of privacy threats in software systems). Then, we introduced how organisational privacy climate [5], i.e., the influence of an organisation's environment on developers' privacy-preserving mindset and behaviour, helps develop privacy-aware software.

A.3 Privacy Enhancing Technologies

The lesson introduced several PETs that can be integrated into software to achieve privacy. Through this lesson, the students realise that privacy can be programmatically achieved in addition to using non-technical methods, such as privacy policies and settings.

> *"PETs are technical measures that allow to utilise personal data while minimising the privacy risks"* [52]

We selected six common PETs discussed in survey studies conducted with software developers: pseudonymisation, k-anonymity, differential privacy, federated learning, homomorphic encryption, zero-knowledge proof, and synthetic data [14, 24, 47, 58, 78]. The lesson introduced the definitions, the applicable scenarios, maturity, weaknesses, and selection criteria of these PETs.

Due to time constraints and to balance student workload, we selected only three PETs for deeper exploration: pseudonymisation, k-anonymity, and differential privacy. We selected these three PETs to provide a balanced perspective on data protection guarantees and the complexity required to learn and implement. Pseudonymisation, representing naive models, is the simplest and easiest to apply. K-anonymity, a probable model, provides stronger guarantees with moderate complexity in learning and applying. Differential privacy, a provable model, is the most advanced of the three regarding learning complexity [31].

A.4 Programming PETs

We taught the students how to program pseudonymisation, k-anonymity, and differential privacy using Python. We selected Jupyter as the programming platform due to its interactiveness and data visualisation properties. During the programming sessions, we also included supplementary questions (e.g.,"K-anonymity: what are the quasi-identifiers in this dataset?, differential privacy: what is the relationship between epsilon and the generated noise?") so that the students gradually learn how to justify the decisions taken when implementing PETs.

Designing programming lessons required significant effort, as only a few PETs-related resources included guidance (e.g., practical exercises, code snippets) to adapt PETs programming. Despite this difficulty, we developed the programming tasks using resources such as [37, 49, 54, 60]. Each programming lesson started with preprocessing datasets and then gradually implementing the specific PET. At any point, students could visualise the outcome in Jupyter (immediate feedback), allowing them to reinforce their learning. The preprocessing step helped students make the dataset compatible with the given scenario and the selected PET. It also made students realise how data preparation assists in privacy protection.

The software scenarios in the questions and datasets were designed using familiar contexts for the students. For example, we

used the dataset "counselling_data", which included 500 dummy data records of university students who participated in counselling sessions. By using such sensitive data, we also wanted to reflect on the importance of protecting privacy. Using the same dataset across different PETs allowed students to observe and compare how each technique processed the data in distinct ways.

Table 7: Learning outcomes aligned with the cognitive levels of the revised Bloom's taxonomy [35].

Cognitive Level	Example Action Verbs	Learning Objective
Remember	recognise, recall, identify	LO1
Understand	explain, summarise, classify	LO2, LO3
Apply	execute, implement	LO7
Analyse	differentiate, organise, attribute	LO5
Evaluate	check, critique, justify	LO6, LO8
Create	generate, plan, produce	LO4

B The Cybersecurity Course

B.1 Prerequisites

To enrol in COMPSCI 316 cybersecurity course, students must have completed 2 prerequisite courses: Computer Organisation and Data Communications and Security. While these prerequisite courses provide essential technical background for understanding computer systems and security, they do not include structured instruction on privacy-related topics.

Computer Organisation. This course introduces the internal workings of a computer system, focusing on how data and algorithms are represented and executed at a low level. Students explore how high-level programming is mapped to the machine level. Key topics include:

- Instruction set architecture
- Data representation
- Assembly (LC3) and C programming
- Memory management

Data Communications and Security. This course provides foundational knowledge of how data travels across networks and the security of data communication. The first part of the course focuses on networking, and the second part introduces security mechanisms to protect data during transmission. Key topics include:

- Network topology
- OSI model
- TCP/IP protocols
- Cryptographic principles and usage

B.2 Lessons Covered Before Week 6

Before introducing the privacy lessons, COMPSCI 316 focuses on core security concepts for understanding system vulnerabilities and protection mechanisms.

- Authentication
- Access control
- Malware
- Risk management
- Secure software development

These lessons do not overlap with the privacy-specific topics introduced in weeks 6 and 7.

C Survey Responses

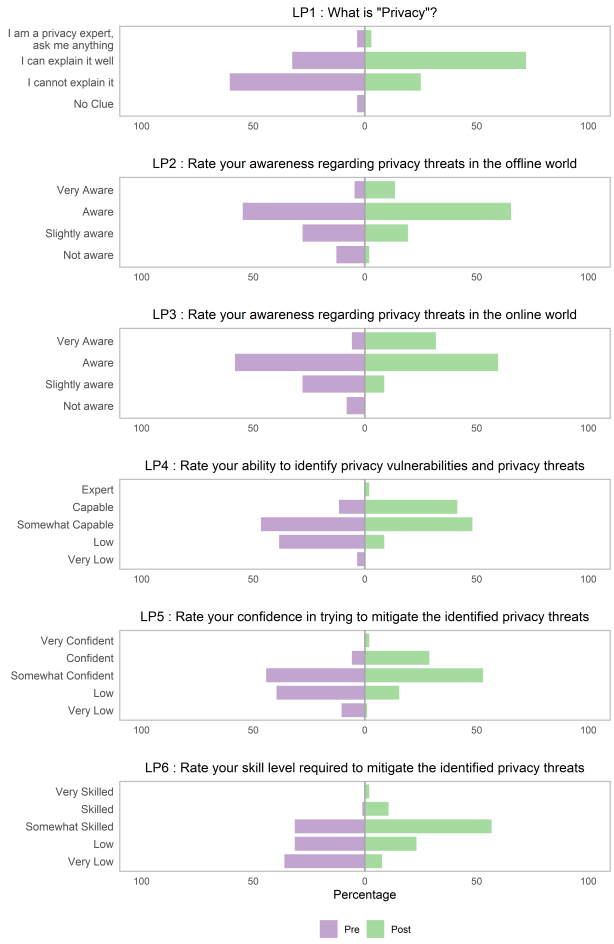

Figure 5: Pre- and post-survey responses of the Likert items LP1 to LP6

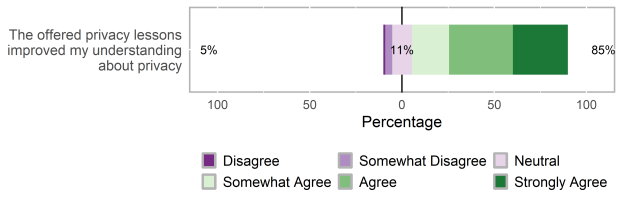

Figure 6: Likert responses for the item LP8

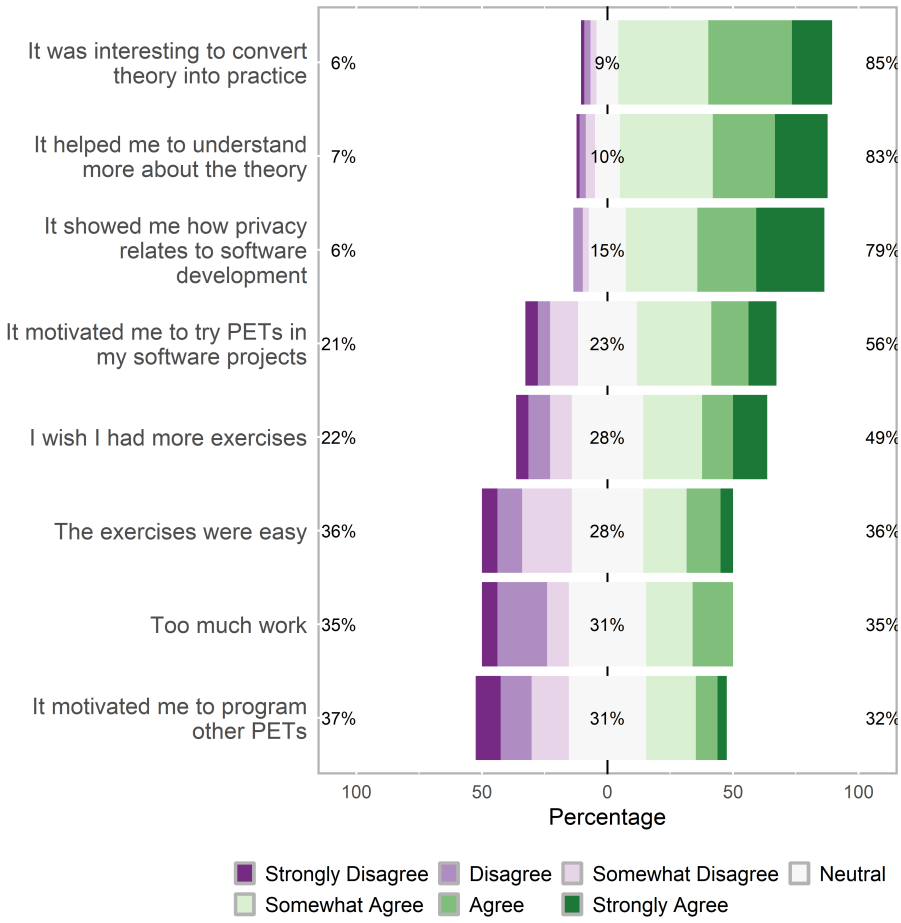

Figure 7: Likert responses for the items PP1 to PP8

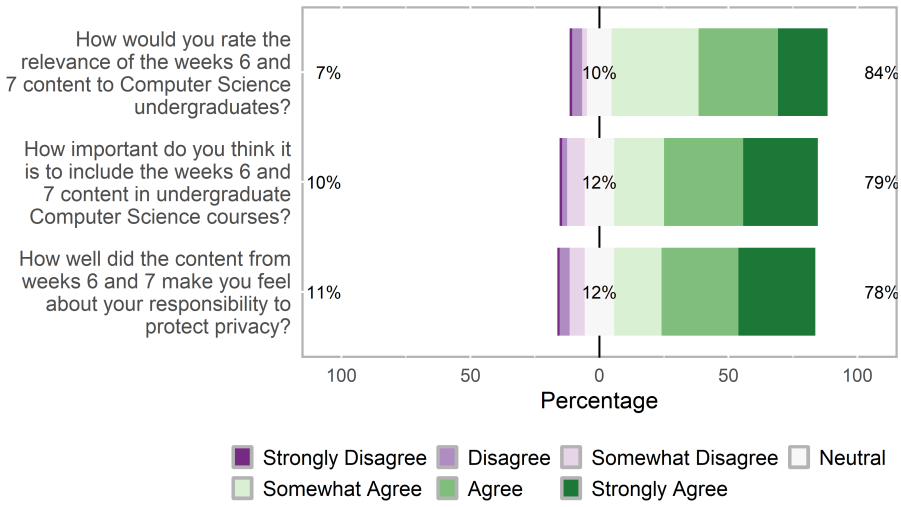

Figure 8: Likert responses for the items PR1 to PR3

D Statistical Tests

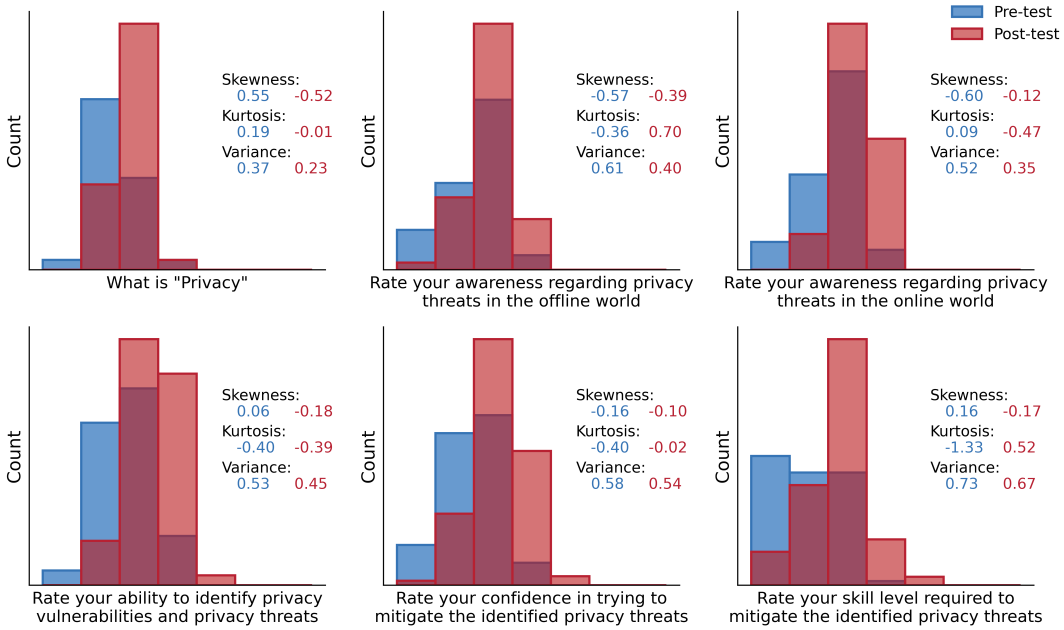

Figure 9: Histograms comparing the pre-survey (blue) and post-survey (red) distributions for the Likert items LP1 to LP6. The skewness, kurtosis and variance show that pre- and post-distributions are non-normal and different in shape.

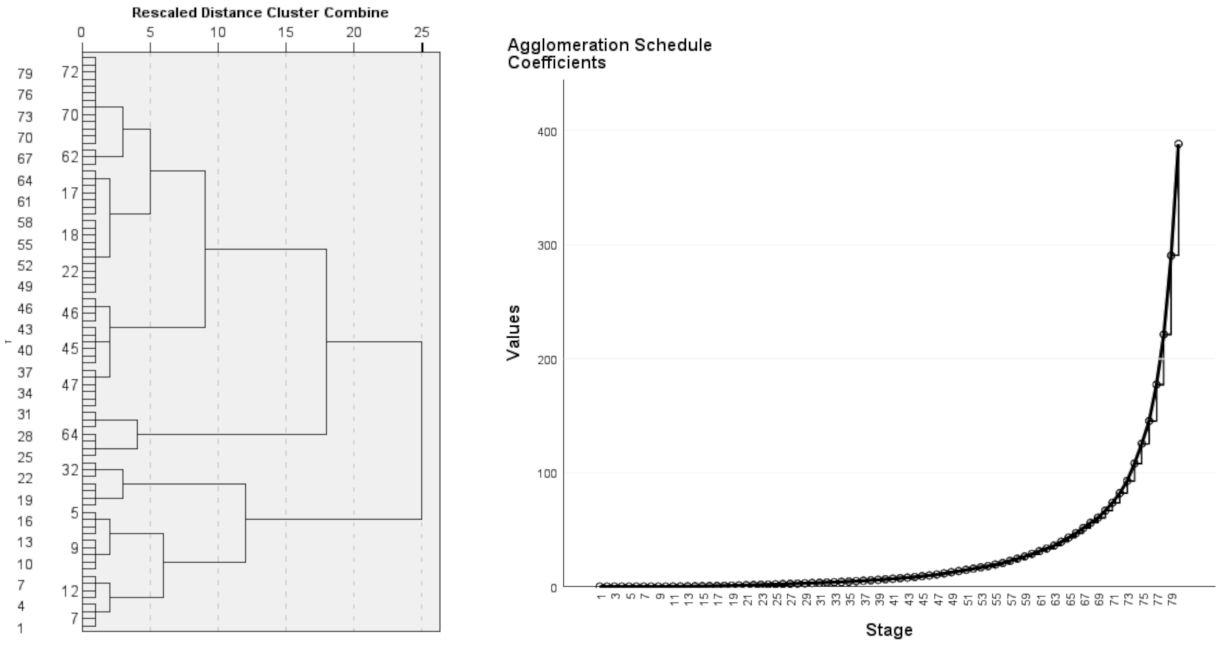

Figure 10: Hierarchical clustering dendrogram (left) using the Ward's linkage. The dendrogram was generated using the SPSS software. The scree plot (right) visually displays how dissimilarity increases at each clustering step.

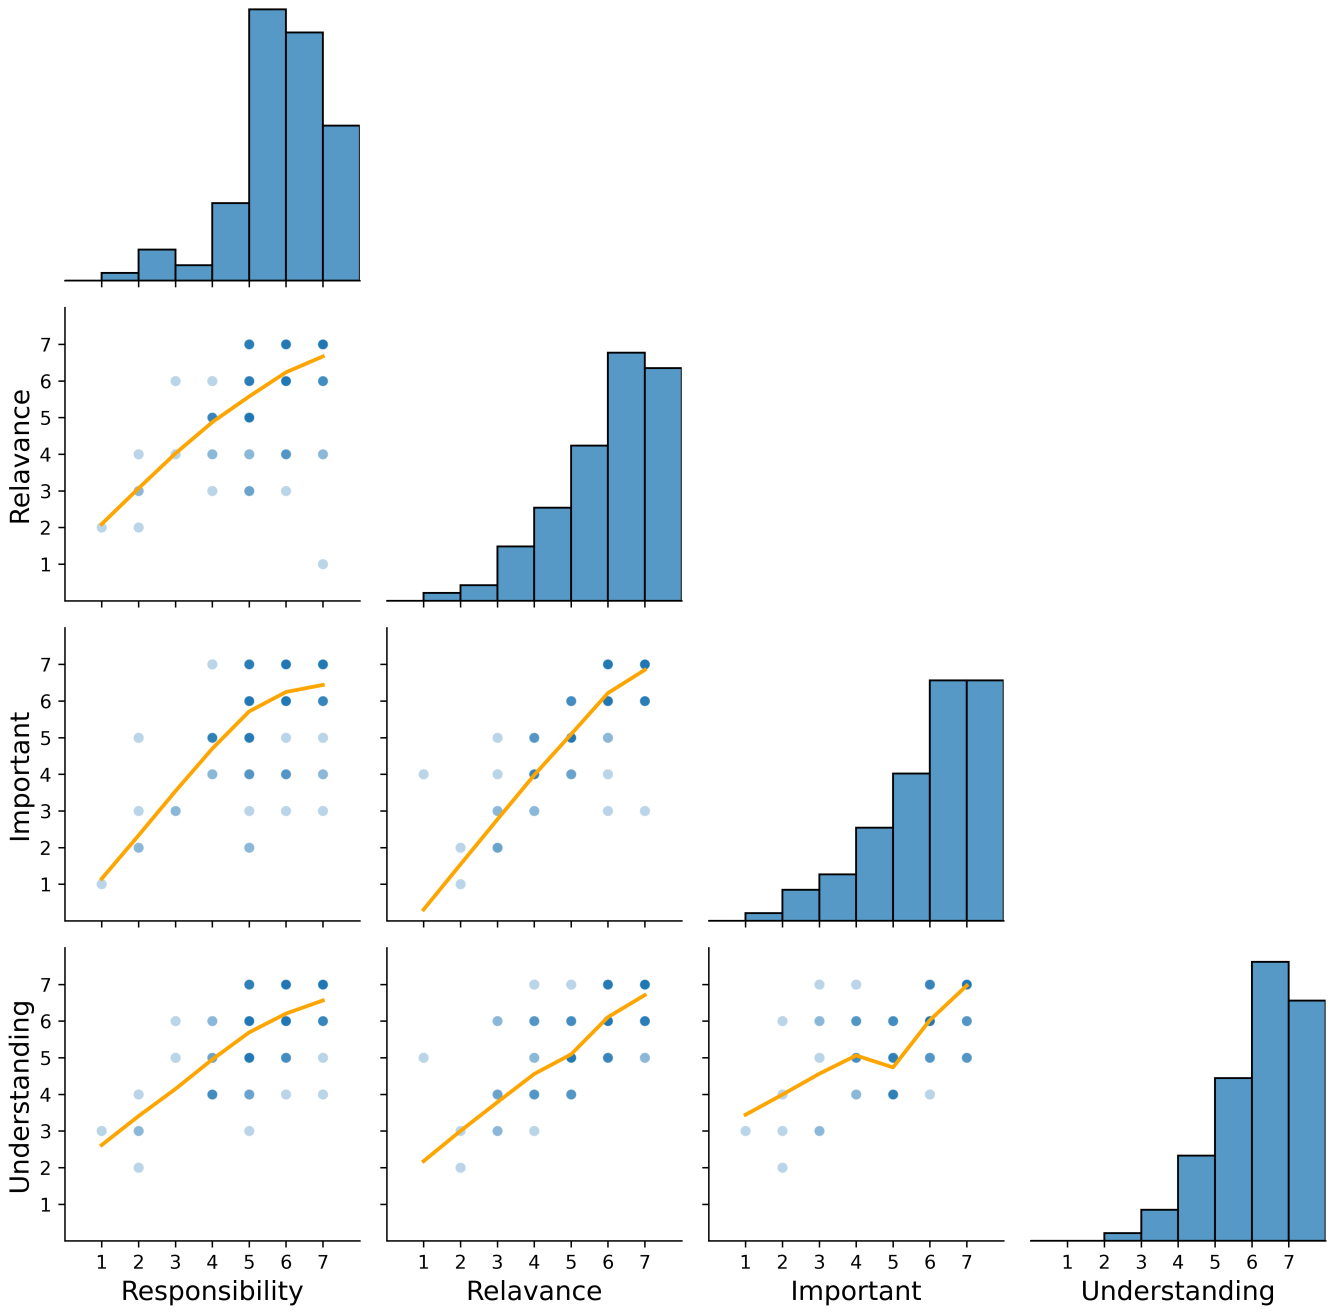

Figure 11: Scatter plot matrix for pairwise relationships between Likert items LP8 (privacy understanding), PR1 (perceived relevance of the privacy lessons), PR2 (perceived importance of the privacy lessons) and PR3 (perceived responsibility in privacy protection). Each scatter plot represents the relationship between two variables, with LOWESS (Locally Weighted Scatterplot Smoothing) curves indicating trends. The transparency of the dots in the scatter plots is inversely related to the frequency of points at a given location. The histograms show that the selected Likert items have a non normal distributions.

E Comparison with Related Work

Table 8: Comparison of the proposed curriculum with related research and privacy-related computing undergraduate courses offered by universities within the top 50 of the Times Higher Education Rankings.

Institute	Intervention / Course Name	Stand-alone Privacy Course	Privacy Lessons					Intervention Evaluation			Artifacts
			General Lessons	Threat Modelling	PETs			Type		Depth	
					Name	Theory	Programming	Quantitative	Qualitative		
University of Auckland	Proposed curriculum	-	•definition •societal aspect •contextual privacy •data privacy •regulations •consequences of privacy breaches-individuals and organisations •privacy by design	✓	Pseudonymisation	✓	✓	✓	✓	Data Collection 1. Pre/post surveys with Likert items, grid selection, and open ended questions 2. Assessments - optional exercises, MST and final exam Results 1. Students improved in privacy knowledge, threat mitigation confidence and skills according to **Mann-Whitney U test, word cloud analysis and median calculation** 2. Identified 5 learner groups based on the perceptions towards programming using **Ward's hierarchical clustering** 3. Using **median calculations**, identified positive trends in perceived role in privacy protection (perceived importance and relevance of the lessons and perceived responsibility to protect privacy) 4. Identified positive correlation between privacy understanding and perceived role in privacy protection using **Kendall's tau-b correlation** 5. Identified 3 themes: early stage of responsibility, privacy-aware software development, and narrowed perspective on positioning privacy knowledge using **Reflexive Thematic Analysis** 6. **Number of attempts** showed that students are motivated by rewards, feedback and different question designs to attempt excercises. 7. **Item Analysis** showed balanced difficulty and discrimination in the mid semester and final exam MCQ questions.	✓
					K-anonymity	✓	✓				
					Differential privacy	✓	✓				
					Federated learning	✓	-				
					Homomorphic Encryption	✓	-				
					Zero-knowledge proof	✓	-				
					Synthetic data	✓	-				
Related Research											
-	[46]	✓	•definition •history •data privacy •web & mobile privacy •adversarial thinking	-	K-anonymity	✓	✓	✓	-	Data Collection 1. Pre/post surveys with Likert items Results 1. According to **Wilcoxon signed-rank test**, student confidence in privacy increased after the course. However, a statistical significant increase was seen only in 5 likert items: debating privacy topics, evaluating ethics and business trade-offs, identifying PETs for personal use and development, and identifying privacy protection techniques for databases.	-
					L-diversity	✓	✓				
					Differential privacy	✓	✓				
University of California, Irvine	[57]	-	•risk assessments •regulations •privacy policies & settings	✓	-			✓	✓	Data Collection 1. Pre/post assessment of multiple choice questions 2. Analysing learning reflection questions Results 1. According to **paired t-test**, students improved knowledge in privacy-social area but not in privacy-technical area 2. Using **Reflexive Thematic Analysis and grounded theory**, lesson feedback were categorised under 4 main themes: learning, course enhancements, application, action	-

✓= provides property; - = missing property; ■ = property not applicable;

Institute	Intervention / Course Name	Stand-alone Privacy Course	Privacy Lessons					Intervention Evaluation			Artifacts
			General Lessons	Threat Modelling	PETs			Type		Depth	
					Name	Theory	Programming	Quantitative	Qualitative		
-	[63]	-	•privacy by design	-	-			✓	-	Data Collection 1. Pre/post surveys - identify ethical issues in 3 scenarios and measure intervention experience Results 1. Did not observe major improvements in identifying ethical issues 2. The interest in learning more about ethical issues in technical work increased according to **frequency analysis** 3. According to **frequency analysis**, students found the intervention interesting and relevant	-
Prairie View A&M University	[38]	-	•location privacy through anonymisation •privacy and utility trade-off	-	-			✓	-	Data Collection 1. Pre/post surveys with Likert items Results According to **mean calculations and frequency analysis**: 1. Students improved awareness, interest, understanding about privacy disclosure, LBS and anonymisation 2. Students thought the lab was effective	-
-	[62]	-	•ethical implications of data collection and use	-	-			-	✓	Data Collection 1. Pre/post surveys 2. Classroom video and audio recording 3. Assessment answers Results 1. Identified 4 themes: novel data experience, empowering vs. unsettling data experiences, experiencing data privacy, and situated knowledge and responsible caring using **grounded theory approach** * Results of the surveys and multimedia data were not provided in this paper	-
University Courses											
University of Oxford	Deep Learning in Healthcare	-	-	-	Federated learning	✓	-				-
					Differential privacy	✓	-				
Harvard University	CS105 Privacy and Technology	✓	•theoretical background •societal aspect •legal perspective •anonymity •re-identification •surveillance •tracing and tapping •emerging technologies (AI)	✓	Differential privacy	✓	-				-
	CS1260 Fairness and Privacy: Perspectives from Law and Probability	✓	•algorithmic foundations for privacy	-	Differential privacy	✓	-				-
Princeton University	COS109 Computers in Our World	-	•personal information •surveillance •tracking •protection measures-users	-	-						✓
	ECE432 Information Security	-	•privacy technologies	-	-						-

✓= provides property; - = missing property; ■ = property not applicable; Course names in blue text are hyperlinks to the respective course pages.

Institute	Intervention / Course Name	Stand-alone Privacy Course	Privacy Lessons					Intervention Evaluation			Artifacts
			General Lessons	Threat Modelling	PETs			Type		Depth	
					Name	Theory	Programming	Quantitative	Qualitative		
Stanford University	CS155 Computer and Network Security	-	•definition •data sharing •tracking •anonymity	-	Tor	✓	-				✓
	CS182 Ethics, Public Policy, and Technological Change	-	•definition •societal aspect •privacy paradox •regulations •anonymisation	-	Differential privacy	✓	-				✓
					Homomorphic encryption (high level)	✓	-				
Caltech	CS162 Data, Algorithms and Society	-	•mentions privacy	-	-						-
University of California, Berkeley	COMPSCI195 Social Implications of Computer Technology	-	•tracking •regulations •threats-users view •protection mechanisms-users	-	-						✓
ETH Zürich	252-0211-00L Information Security	-	•motivation and definitions •policies and policy languages •mechanisms •anonymity •application case studies: mix networks and crowds	-	-						-
University of Chicago	CMSC10434 Technology and Privacy in the Digital Age	✓	•historical foundations •societal aspect •cultural aspect •policies	-	-						-
	CMSC23206 Security, Privacy, and Consumer Protection	-	•consumer privacy •censorship •platform content moderation •data breaches •government surveillance •regulations	-	-						-
	CMSC23210 Usable Security and Privacy	-	•regulations •privacy notices •anonymity •online data collection	-	-						✓
	CMSC23218 Surveillance Aesthetics: Provocations About Privacy and Security in the Digital Age	-	•contextual integrity •anonymity •privacy notices •data-driven privacy tools •user-perspective	-	-						-
	CMSC23800 Adversarial Machine Learning	-	•privacy of ML models •privacy attacks	-	-						-
	DATA25900 Ethics, Fairness, Responsibility, and Privacy in Data Science	-	•privacy issues	-	-						-
	CMSC25910 Engineering for Ethics, Privacy, and Fairness in Computer Systems	-	•privacy invasiveness of computer systems •algorithmic approach	-	-						-

✓= provides property; - = missing property; ■ = property not applicable; Course names in blue text are hyperlinks to the respective course pages.

Institute	Intervention / Course Name	Stand-alone Privacy Course	Privacy Lessons					Intervention Evaluation			Artifacts
			General Lessons	Threat Modelling	PETs			Type		Depth	
					Name	Theory	Programming	Quantitative	Qualitative		
Johns Hopkins University	EN.601.104 Computer Ethics	-	•privacy issues	-	-						-
	EN.601.443 Security & Privacy in Computing	-	•mentions privacy	-	-						-
	EN.601.124 The Ethics of Artificial Intelligence and Automation	-	•mentions privacy	-	-						-
National University of Singapore	CS4267 Algorithmic Foundations of Privacy	✓	•anonymity •data privacy •privacy attacks (inference & reconstruction)	-	Differential privacy	✓	-				-
					Tor	✓	-				
University College London	COMP0061 Privacy Enhancing Technologies	✓	•private communications •anonymous communications •traffic analysis •interdisciplinary aspects •cryptographic protections	-	Differential privacy	✓	-				
					Zero-knowledge proof	✓					
	COMP0056 People and Security	-	•data protection •privacy by design •PST model •Surveillance, dataveillance, and sousveillance	-	-						-
Carnegie Mellon University	15330 Introduction to Computer Security	-	•mentions privacy	-	-						-
	15316 Software Foundations of Security & Privacy	-	-	-	Differential privacy	✓	-				✓
Duke University	COMPSCI351 Computer Security	-	•technologies to support online privacy (not listed)	-	-						-
Northwestern University	COMPSCI496 Security and Privacy Education	-	•analysing privacy education approaches for users and technology designers (lesson plan not given)	-	-						-
	COMPSCI312,412 Data Privacy	✓	•data privacy •database anonymisation •anonymous communications •algorithmic fairness •privacy in web, social media and ML	-	Differential privacy	✓	-				-
	COMPSCI396: Differential Privacy: from Foundations to Machine Learning	✓	•data privacy •privacy attacks •algorithms for private learning	-	Differential privacy	✓	-				-

✓= provides property; - = missing property; ■ = property not applicable; Course names in blue text are hyperlinks to the respective course pages.

Institute	Intervention / Course Name	Stand-alone Privacy Course	Privacy Lessons					Intervention Evaluation		Artifacts
			General Lessons	Threat Modelling	PETs			Type	Depth	
					Name	Theory	Programming	Quantitative / Qualitative		
École Polytechnique Fédérale de Lausanne	COM301 Computer security and privacy	-	•mentions privacy	-	-					-
Georgia Institute of Technology	CS4726 Privacy Tech Policy	✓	•privacy in technology, policy, ethics, law, and business	-	-					-
University of British Columbia	COSC_O421 Network Science	-	•data privacy	-	-					-
McGill University	COMP189 Computers and Society	-	•data privacy	-	-					-
	COMP555 Information Privacy	✓	•privacy by design •privacy threats •privacy concerns in databases, web, mobile apps and cloud	-	-					-
University of Illinois at Urbana-Champaign	CS211 Ethical and Professional Conduct	-	•mentions privacy	-	-					-
	CS442 Trustworthy Machine Learning	-	•membership and model inversion attacks •differentially private data generative models	-	Differential privacy	✓	-			-
	CS461 Computer Security I	-	•assess and address privacy issues for policy and humans •privacy risk analysis according to CIA triad •human issues in privacy •legal and ethical issues	-	-					-
	CS463 Computer Security II	-	•privacy & anonymity •policy composition and analysis •privacy of emerging systems •privacy issues in social networks •privacy issues in web •human factors in privacy	-	-					-
	CS464 Topics in Societal and Ethical Impacts of Computer Technology	-	•mentions privacy	-	-					-

✓= provides property; - = missing property; ■ = property not applicable; Course names in blue text are hyperlinks to the respective course pages.

F Codebook

Table 9: Codebook of the Reflexive Thematic Analysis, with associated themes, sub-themes, and codes.

Theme	Sub-th.	Codes - Round 2	Codes - Round 1	Example Comments
Early Stage of Responsibility		Vigilant in privacy threats	Vigilant in privacy threats	Now I have a more baseline understanding of privacy importance, and in particular ways it can be threatened.
		Vigilant in their own privacy	Vigilant in privacy threats	It's possible that it could make me more aware of protecting my privacy and the privacy of the content of my work.
		Vigilant in their own and others' privacy	Vigilant in their own privacy	It can make me pay more attention to privacy protection in my future work, both for myself and for users.
		Vigilant in their own privacy	Vigilant in their own privacy	It will give me the fundamental knowledge. It will make me more aware of preserving privacy in practice.
		Vigilant in their own privacy	Privacy protection as a user	Be aware of sharing my personal details through everything I interact with.
		Vigilant in their own privacy	Vigilant in their own privacy	My understanding of all the content seems like it could be potentially relevant. Understanding what privacy/personal data is, the threats to it, and potential ways to protect it is useful information.
		Vigilant in protecting privacy	Vigilant in protecting privacy	I plan to work with technology in the future, so privacy will be very important and information/data will need to be protected.
		Vigilant in their own privacy	Vigilant in their own privacy	It will help me protect my personal information in the future and could be useful in the industry I would like to work with.
		Vigilant in their own privacy	Privacy protection as a user	It helps improve my awareness of privacy in my study. When I search for some information on the internet, I will be more careful to share my information.
		Vigilant in their own privacy	Privacy protection as a user	Be careful about your personal data and protect your privacy.
		Vigilant in protecting privacy	Considering to protect user data	Helped develop a stronger understanding of protecting users' data.
		Compliance obligation	Compliance	Good to know for compliance.
		Compliance obligation	Compliance	Understanding data privacy protection regulations and how to implement them.
		Compliance obligation	Privacy protection motivated by regulatory consequences	Privacy security has its own set of laws that can result in big consequences if companies do not comply. Therefore, I think it is very important to know how to implement these through code.
Privacy-Aware Software Development	Contribution through Programming	Understanding privacy can be integrated into programmers' workflow	Protect privacy through programming	Gives me a better idea of how privacy is used in programming.
		Understanding privacy can be integrated into programmers' workflow	Intention to protect privacy through programming	Assignment 2 made me more familiar with the PETs we covered, and the experience implementing them will be good to draw from.
		Understanding privacy can be integrated into programmers' workflow	Understanding the usefulness of privacy programming	Helped me see how the techniques are implemented.
		Understanding privacy can be integrated into programmers' workflow	Understanding the usefulness of privacy programming	Before weeks 6 & 7, I had trouble understanding how our knowledge in cybersecurity could be put into practice in the real world, and getting a glimpse of it through coding different PETs was very helpful in that understanding.
		Understanding their skills can be beneficial to protect privacy through coding	Protect privacy through programming	Helped in knowledge and application of PETs to prevent privacy vulnerabilities where they might not be seen.
		See privacy programming as a professional skill	Protect privacy through programming	I have tried to code for privacy, which did make me think about how real professionals achieve that. So that can be useful for my future career.
		See privacy programming as a professional skill	Protect privacy through programming	An understanding of how privacy can be implemented in an application is helpful because it should be a key requirement for any developer.

Theme	Sub-th.	Codes - 2nd Round	Codes - Round 1	Example Comments
Privacy-Aware Software Development	Privacy-First Approach	Privacy-first software development	Privacy-first software development	I will now consider and implement privacy mechanisms during the programming/creation stage of a project.
		Privacy-first software development	Privacy-first software development	I'll definitely approach software development with a privacy-first mindset.
		Privacy-first software development	Privacy-first software development	It allowed me to notice more privacy issues in the world around me, and I think it will be helpful in the future so that I can incorporate privacy principles into my designs from the beginning stages.
		Privacy-first software development and disagreement with ignoring privacy	Privacy-first software development	Being more conscious of building privacy protection measures into software projects rather than as an addition.
		Privacy-first software development and disagreement with ignoring privacy	Privacy-first software development	I'm genuinely surprised that privacy isn't a permanent topic in COMPSCI 316. I think this will definitely help me prioritize privacy in future projects as it is already often forgotten and only implemented when there is a privacy breach.
	Ownership in Privacy Protection	Internal commitment to protect privacy	Sense of agency in protecting privacy through coding	I will be implementing a PET (K-anonymity) into my own personal project.
		Desire to improve privacy knowledge by self-learning	Desire to improve privacy knowledge by self-learning	Basic understanding of privacy so I can search for more to learn more to implement in work or even further studies.
		Awareness of privacy flaws in past work and readiness to improve	Awareness of privacy flaws in past work and readiness to improve	It tells me the importance of protecting privacy. I know why my previous website work is unsafe and know how to improve it.
		Internal commitment to protect privacy	Internal commitment to protect privacy	As someone who currently works in one of the largest tech companies, we have a lot of policies regarding data privacy protection. I was following them out of compliance, knowing that violations lead to paperwork. The lecture content regarding data privacy helped me understand why we need to protect data privacy in the first place. It has also given me insight into why I have to complete our mandatory training and helped me understand why there is a need to put these training modules into practice. Learning about the impact of data privacy protection helped me be more proactive in protecting our customers' data and helped me understand our internal processes regarding data handling.
Narrowed Perspective in Positioning Privacy Knowledge		Uncertainty in using privacy knowledge	Uncertainty in using privacy knowledge	Unsure.
		Uncertainty in using privacy knowledge	Uncertainty in using privacy knowledge	No clue.
		Uncertainty in using privacy knowledge	Uncertainty in using privacy knowledge	Not much, not that relevant for what I want to do.
		Uncertainty in using privacy knowledge	Uncertainty in using privacy knowledge	I'm not sure if it will.
		Uncertainty in using privacy knowledge	Uncertainty in using privacy knowledge	I don't think it will be particularly useful for me.
		Reluctance to put privacy into practice	Difficulty in putting privacy into practice	Not too much, it improved my understanding but putting it into practice looks too much of a hassle.
		Limited understanding of how to use the privacy knowledge	Limited understanding of how to use the privacy knowledge	It was really good to learn about the few common PETs, and I believe that it'll be useful in any form of data handling. It's quite rare for us to think of what will happen in the future and if we're going to be implementing any of the PETs learned in class, although it's good to have an understanding.
		Inability to map the privacy knowledge in diverse professions	Limited understanding of how to use the privacy knowledge	If I end up working in the field of cybersecurity, it will be very helpful, including that I will be able to communicate with people more calmly in future employment and will not offend them.
		Inability to map the privacy knowledge in diverse professions	Limited understanding of how to use the privacy knowledge	Basically, any IT industry will require privacy knowledge.