Peer Instruction for Digital Forensics

William Johnson, Irfan Ahmed[*], Vassil Roussev
Department of Computer Science
University of New Orleans
wejohns1@uno.edu, {irfan,vassil}@cs.uno.edu

Cynthia B. Lee
Computer Science Department
Stanford University
cbl@stanford.edu

Abstract

Digital forensics can be a difficult discipline to teach effectively because of its interdisciplinary nature, closely integrating law and computer science. Prior research in Physics and Computer Science has shown that the traditional lecture approach is inadequate for the task of provoking students' thought-processes and systematically engaging them in problem-solving during class. Peer instruction is an established pedagogy for addressing some of the challenges of traditional lectures. For this paper, we developed 108 peer instruction questions for a digital forensics curriculum, and evaluated a selection of the questions by holding a condensed computer forensics workshop for university students. The evaluation results show that peer instruction helps students understand the targeted digital forensics concepts, and that 91% of students would recommend that other instructors use peer instruction.

1 Introduction

Digital forensics is defined as the application of scientific tools and methods to identify, collect, and analyze digital artifacts in support of legal proceedings [11]. It is a challenging field to teach effectively because of its positioning at the intersection of law and computer science. For instance, the data acquired from a suspect computer (a.k.a digital evidence) without a search warrant may prevent the evidence from being admitted in court.

Students need to achieve a clear understanding of important principles of performing a forensic investigation within the constraints put forth by state and federal law. Often, forensic instructors combine traditional lectures with hands-on lab exercises to improve the student learning outcomes. Typical exercises include solving cybercrime cases such as theft of credit card data and malware attacks. In our experience, forensic hands-on exercises are effective for students who have a good understanding of individual forensic concepts relating to the lab task, and are able to integrate them into a broader view of a forensic investigative scenario.

Unfortunately, the traditional lecture-centric approach is often inadequate in providing the conceptual preparation for hands-on labs and exercises. Commonly, lectures fail to motivate students to study before class, and therefore the in-class learning experience is less effective than it could be. The traditional lecture also lacks a reliable in-class mechanism to help an instructor ensure that students understand the concepts being taught. Although the instructor may ask ad-hoc questions to stimulate discussion, this usually engages a few students and does not provide feedback about every student in class.

Peer instruction has emerged as a promising solution to enrich the in-class experience such that it yields better outcomes. It was first introduced by Eric Mazur in physics classrooms at Harvard University [2]. With peer instruction, a lecture consists of a series of multiple-choice questions (a.k.a. ConcepTests) that are designed to provoke deep conceptual thinking in students and engage them in a meaningful discussion. Students are required to study some reading material before coming to class, which helps them to find correct answers to the peer instruction questions. Peer instruction has shown to be effective in physics, computer science, and biology classrooms [7]. In particular in computer science, it has halved failure rates in four courses [7] and increased student retention in the computer science major [10].

Previously, peer instruction has not been explored as a means for delivering a digital forensics curriculum. We have developed 108 peer instruction questions for an introductory course on computer forensics. This paper discusses some of these questions as examples and analyzes classifies with respect to two basic criteria. First, what are the elements used in the questions to trigger conceptual thinking processes of students such as "compare and contrast," "deliberate ambiguity" and "trolling

[*]Corresponding author

This work was supported by the NSF grant # 1500101

for misconception"? We use Beatty et al.'s concept triggers [1] for the analysis. Second, how are the questions presented, such as in a scenario, example, or diagram.

The paper also presents the results of a digital forensic workshop utilizing peer instruction methodology and questions. The evaluation results show that most of the students find the use of clickers beneficial and the group discussion for a peer question helps the students understand a target concept. 91% students would recommend other instructors use peer instruction. The results also show a clear evidence of improvement in student learning that is measured through *quizzes* taken before and after a topic is covered and *peer instruction questions* replied to by the students before and after peer discussions.

The rest of the discussion is organized as follows. Section 2 discusses the peer instruction methodology and related work. Section 3 presents the elements of a peer instruction question (i.e. concept triggers and question presentation) along with the steps to create a peer question. Section 4 discusses four examples of peer instruction questions and identifies their concept triggers and question presentation types, followed by section 5 that presents an analysis of 108 peer instruction questions recently developed by the authors. Sections 6, 7, and 8 present the peer instruction implementation in a digital forensic workshop, a list of peer instruction questions used in the workshop, and evaluation results, followed by conclusion in section 9.

2 Background/Related Work

2.1 Peer Instruction Methodology

Peer instruction pedagogy could be categorized under the general umbrella of a"flipped classroom" approach, but the term "peer instruction" describes a specific protocol, and does not mean generally any use of peer discussion or clickers. The protocol has three important elements. First, it divides a class lecture into a series of multiple-choice conceptual questions. The questions focus on the primary concepts an instructor wishes to teach. Second, it divides the students in class into a number of small groups for discussions. Third, it requires students to read some material before coming to class. The material is related to the topic covered in the class and enables the students to find answers of the questions.

For each question, the following steps are involved:

- A peer instruction question is posed to students. The students use clickers to reply.
- The instructor is able to view the results. If the answers are mostly incorrect, the instructor then encourages the students to discuss their answer choices with their group members.
- The instructor then poses the same question to the students again; the students reply to the question.
- Depending on the answers, the instructor would have the option to discuss the concept in further detail with the students to ensure a better understanding; or, if the students clearly understand the concept at this point, the instructor may choose to continue to the next question.

2.2 Related Work

Peer instruction has not been widely adopted in the cybersecurity classroom. Recently, Johnson *et al.* [4] developed peer instruction questions for two courses: introduction to computer security, and network penetration testing. They developed a set of 172 peer instruction questions; however, these were not evaluated in a classroom setting.

More broadly, there have been a number of prior efforts to use peer instruction in the computer science classroom. Porter *et al.* [8] perform a multi-institutional study of the usage of peer instruction in seven instructors' introductory programming courses. Considering instructors' prior experience (or lack thereof) utilizing peer instruction, they primarily focus on student perception of peer instruction using measurements such as perceived question difficulty, question time allowed, discussion time allowed, and content difficulty. They obtain participants' responses using surveys and note that at least 71% of students would recommend other instructors to use peer instruction.

Similarly, Porter *et al.* conduct a measurement of peer instruction across multiple small liberal arts colleges to measure the effectiveness of peer instruction in smaller classes, using data from five instructors at three institutions [9]. They notice normalized gains in the same range that of larger universities with students generally approving of the method and their performances.

Esper [3] utilizes peer instruction in a software engineering course that has 189 students. She makes a modification in the standard peer instruction process in which a clicker question is initially shown without answers and both the students and instructor propose potential answer choices with discussions of those answers. It is worth noting that the instructor does not mention whether an answer suggestion is correct or incorrect. The evaluation results show that 72% of the students would recommend the course instructor, with 28% not recommending due to the reasons such as correct answers are not given and clicker questions are unclear.

Liao *et al.* [6] create models of in-class clicker questions to predict low-performing students early in the term. They use a linear regression model to predict final exam scores with approximately 70% accuracy in a twelve-week introductory computer science course.

Lee *et al.* [5] examine the effectiveness of peer instruction in two upper-level computer science courses: Theory of Computation and Computer Architecture and find the average normalized learning gains of 39%.

3 Developing Peer Instruction Questions

A peer instruction question has two important elements: concept trigger(s) and question presentation. A concept-trigger is a hidden element in multiple choices of a question that is deliberately introduced in the question. The hidden element triggers the students thought process and helps them with useful peer discussion allowing them to understand the target concept and further eliminate either incorrect choices or identify a correct choice. We utilize concept-triggers from Beatty *et al.* [1]. Table 1 presents a list of concept-triggers used in our questions such as *Use "none of the above"*, *Omit necessary information*, and *Trap unjustified assumptions*.

A question presentation refers to how a question is written or articulated. We use five presentation types: scenario, example, definitional, diagram, and feature. A concept may be comprised of multiple components or features. If a question or its multiple choices target these features, we classify the question as *feature*. The question asks students to identify required features of a concept in multiple choices, or a feature is provided in a question which asks students to identify its corresponding concept in the choices.

To develop a question, an instructor should first thoughtfully select a target concept. Peer instruction questions do take a significant amount of class time, and so questions should focus on the most essential elements of the course. Second, the instructor should select a question presentation that is suitable to articulate the concept. Finally, the instructor should choose at least one concept trigger. A question can have multiple concept triggers (examples are given in the next section).

4 Examples of Peer Instruction Questions

4.1 FAT32 File System Internals

In the FAT32 file system, there are two primary data structures used to maintain knowledge of data: the directory entry and the file allocation table. A directory entry is present for every file and folder and contains the file name, file size, the address of the starting cluster of a file's content. The file allocation table is used to identify the next cluster in a file and the allocation status of clusters. The final cluster of a file has a marking `EOF` in its FAT entry that indicates end of file. Therefore, in order to access a specific file from a FAT file system, the filename and starting cluster are necessary.

A peer instruction question on this concept can be created as follows: *In order to access a file from a FAT file system, what information is absolutely necessary?* a) Name and ending address of file content; b) Name, file size, and ending address of file content; c) Name and starting address of file content, d) File size and starting address of file content, e) None of the above (answer: c).

Concept trigger. We use "none of the above" to present a possibility that the provided choices may be incorrect. We also use "trap unjustified assumptions" by introducing file-size in choices–FAT file system's data structures allow file clusters to be accessed similarly to a linked-list, whereas the last cluster is marked as `EOF` thereby does not require file size to access a file.

Question presentation. This question is presented as a feature question as it draws upon and asks about core features of the FAT file system and its data structures.

4.2 File Carving

File carving is an important concept in computer forensics. It is a method used to recover files when they are inaccessible through file system, such as deleted files, or a corrupted file system. In many cases, the file content on disk remains intact and can be recovered by file carving.

A basic form of file carving assumes that a file blocks are physically located in a sequence on disk and requires only the location of the first and last block of a file to recover the file content. It searches for file headers and footers using predefined signatures to identify the first and last blocks of files. This technique however, suffers in the presence file fragmentation that can cause a file's blocks to be store non-contiguously.

In addition, while defragmenting a drive places allocated file blocks in contiguous pieces, the defragmentation process may overwrite blocks that have been marked by the file system as deallocated.

A peer instruction question on basic file carving addressing the above-stated challenges can be created as follows: *In which of the following situations is file carving most effective?* a) The targeted drive is highly fragmented, b) The targeted drive has been recently defragmented, c) The system being used to examine the drive has low free space, d) The system being used to examine the drive has high free space, e) More than one of the above (answer: d).

Concept trigger. Deconstructing the question, we note that two triggers are used. First, as we provide the potential for multiple question choices in option (d), we use "identify a set or subset."

To further assess a student's understanding of the downsides of both fragmentation and defragmentation processes, we use "trolling for misconceptions", as option (b). The defragmentation presents an ideal layout of contiguous data (and thus is an attractive option), how-

Concept-trigger Name	Brief Description
Compare and contrast	Compare multiple situations; draw conclusions from comparison
Interpret representations	Provide a situation that asks students to make inferences based upon the presented features
Identify a set or subset	Ask them to identify a set or subset fulfilling some criterion
Strategize only	Provide a problem; ask students to identify the best means of reaching a solution
Omit necessary information	Provide less information than is essential for answer; see if students realize this
Use "none of the above"	An option to learn alternative understandings; use it occasionally as a correct answer
Qualitative questions	Questions are about the concepts and relationships rather than numbers or equations
Analysis and reasoning questions	Questions require significant decision-making, hence promote significant discussion
Trap unjustified assumptions	Answer choices are facilitated by potential unjustified assumptions made by the students
Deliberate ambiguity	Use deliberate ambiguity in questions to facilitate discussion
Trolling for misconceptions	Attempt to trap students with answers that require common misconceptions to choose

Table 1: List of concept-triggers from Beatty *et al.* [1]

ever, it also raises a possibility that important data has been lost.

Question presentation. This question does not specifically present a scenario, but provides examples of potential carving situations on both a target drive and an investigator's machine. This question thus uses the "example" presentation type.

4.3 Registry Forensics

Forensics of the Microsoft Windows registry provides significant details of activities on a computer. The registry is used by the operating system and applications as a database for storing configuration information. It consists of *hives* (backed by files) and can hold information such as list of applications to run on user login, user password hashes, the last time particular USB devices were plugged into the system, and the list of installed software. In a registry hive, *keys* are similar to directory paths, *values* are analogous to file names, and *data* is analogous to file content. The *SYSTEM* hive has a key called *USBSTOR* that maintains a set of USB storage devices plugged into a system, holding their serial numbers with the last time the devices were inserted into the system.

A peer instruction question on registry can be created as follows: *A USB drive with an unknown owner is found in a corporate setting. How might a forensic investigator typically determine whether that particular drive was plugged into any given Windows machine?* a) Examine all *ntuser.dat* files to determine if a user plugged it into the machine; b) Check the *SYSTEM* registry hive to see if it was plugged into that machine; c) Check the SOFTWARE registry hive to see if it has been used by any particular piece of software; d) More than one of the above; e) None of the above (Answer: b).

Concept trigger. The question utilizes "use *none of the above*" to point the students toward the possibility of each answer being potentially invalid. However, with multiple potentially valid options, it also uses "identify a set or subset" as option d). While each option is possible and may assist in locating references to USB drives, the use of the SYSTEM hive is the most efficient. SYSTEM hive contains system-wide activities. If no reference to the drive is found in the hive, there is no need to search through installed software or any user-specific locations.

Question presentation. This is a scenario question. It provides a particular situation that a forensic investigator may face, and it asks for the best means to proceed.

4.4 Forensic Artifacts

Redaction of information is done generally through obscuration (with pens, opaque tape, etc.) or excision of the document through cutting the intended area out of the document. Improper redaction, like many other processes (such as file deletion), can create unintended forensic artifacts.

A peer instruction question on this concept can be created as follows: *Which of the following are examples of redaction?* a) A company employee selectively covers sentences from a press release with a black marker; b) A copy editor at a nonprofit adds information to a press release draft with a red pen; c) An inside attacker at a law firm cuts documents to physically remove appearances of his name; d) Both a) and b); e) Both a) and c) (answer: e).

The question checks whether students can readily identify examples that illustrate the core features of redaction, i.e., actions that obscure, or remove, information.

Concept trigger. As the question is somewhat simple and looking for provided examples of the essential features of redaction, it uses "analysis and reasoning" as students scan the answers for information that suggests obfuscation or removal.

Question presentation. The question requires students to demonstrate the understanding of the features of the Redaction concept; it is therefore a "feature" question.

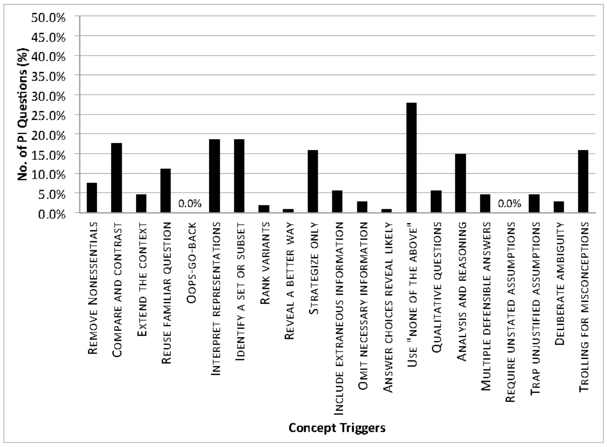

Figure 1: Percentage of peer instruction questions (of computer forensics) over concept triggers

Topics	# of Questions
Introduction to Computer Forensics	31
Windows Registry	10
Forensic Artifacts	24
File Systems	11
Live Forensics	24
File Carving	8
TOTAL	108

Table 2: Number of peer instruction question developed for different computer forensics topics

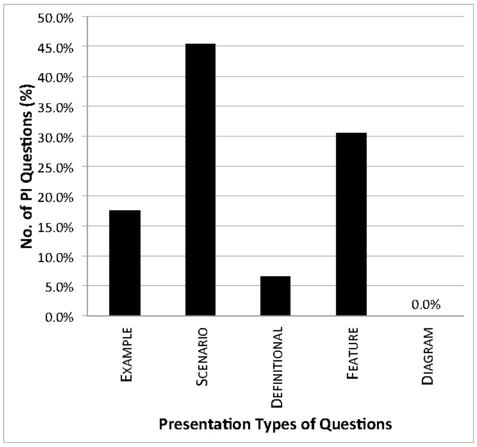

Figure 2: Percentage of peer instruction questions (of computer forensics) over presentation types

5 Analysis of the Forensics Questions

We developed 108 peer instruction questions for the computer forensics course covering six topics (refer to Table 2). Due to the limited space, we cannot present all the questions with their detailed analyses as presented in the last section for four example questions. This section presents an analysis of questions in the context of concept-triggers and presentation types to provide some insight.

5.1 Concept Trigger

Out of the 108 questions, 67 questions use multiple concept triggers. If we exclude "use *none of the above*", as this has been used commonly during question development, there are still 48 questions with multiple triggers.

Figure 1 presents the percentages of questions that utilize each concept trigger. A number of concept triggers are conspicuously missing from the graph–we attribute this to the instructor's personal preferences as well as difficulty of utilizing some concept triggers, such as "oops-go-back", as some of these require a lecture to be in perhaps greater detail than desired to utilize sequences of closely-related peer instruction questions. The concept triggers are considered commonly used if they are used more than 15% on average. They include "compare and contrast", "interpret representations", "identify a set or subset", "Use *none of the above*", "analysis and reasoning", and "trolling for misconceptions".

5.2 Question Presentation Types

Figure 2 shows the distribution of peer instruction questions over question presentation types. Almost 45% questions are scenario based and 18% questions are based on examples. These two types of questions are naturally aligned with the forensic courses' content as they are often based on case studies and examples. The diagram questions are missing because they are time-consuming and difficult to develop. It is worth mentioning that the developed questions reflect the authors' preferences.

Figure 3 further presents the distribution of peer instruction questions over presentation types for different forensic topics. "Windows forensics" and "file systems" have mostly feature-based questions. Scenario-based questions are mostly present in introduction to computer forensics, and live forensics topics. The "forensic artifacts" topic utilizes scenario, example, definitional, and feature types equally. "Definitional" type of questions are mostly used in file carving.

6 Peer Instruction Implementation

To evaluate the effectiveness of peer instruction in computer forensics, we implement peer instruction methodology in a four-hour workshop on digital forensics at the University of New Orleans covering four topics, introduction to computer forensics, file carving, Windows registry, and FAT file system.

Workshop Advertisement and Student Participants.

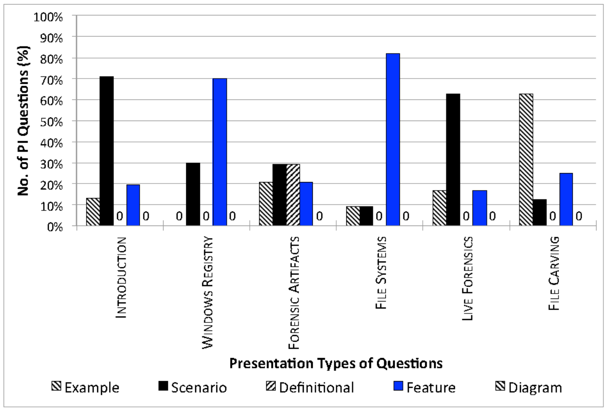

Figure 3: Percent distribution of computer forensics questions as per presentation types and topics

We advertised the workshop approximately 10 days prior through an email to students. The email contains a brief overview of computer forensic investigation and peer instruction. It also has a link to a signup form, which has name and email address fields, and further inquires the students about their willingness to read some material before the workshop. The form also asks them whether they have taken the following courses to assess their knowledge-level: data structures, introduction to computer forensics, principles of operating systems, or any other cybersecurity courses. We received 20 responses to the interest form and ultimately 12 attended the workshop, which is meaningful as a minimum graduate and undergraduate classes require 7 and 14 students respectively at the University of New Orleans.

Pre-class Reading Material. We provided two pieces of reading material. The first was a short paper detailing the importance, history, structure of, and useful data found in the Windows registry. The second consisted of the first five pages of an article introducing file carving, touching on the FAT32 and NTFS file systems as well as their file allocation and deletion procedures, detailing fragmentation and means of file recovery in both, as well as carvers utilizing header and footer file magic.

In addition, we distributed a quiz of five questions using Google Forms to the respondents of the interest form. If the respondents complete the quiz, it confirms their interest of attending the workshop and also better ensures that they read the material and prepare for the workshop. The quiz questions are available in appendix B.

For the quiz, we received 11 responses. 7 of those answered all questions correctly, and 4 respondents answered with 3/4 questions correct.

In-class Peer Instruction Activities. The twelve students were divided into four groups based-on their proximity i.e. one group of four students, two groups of three students and one group of two students. Seating was restricted to the front of the room where clickers were placed with workstations. The lecture was centered around seven peer instruction questions, distributed among four separate lecture sections: "Introduction to Computer Forensics", "File Systems", "File Carving", and "Windows Registry". As we began each question, we provided the students 2-3 minutes to respond. 3-5 minutes were provided for them to discuss their answers among themselves, and we provided another 2-3 minutes for the second set of responses to the question.

7 Workshop Peer Instruction Questions

This section lists the seven peer instruction questions used in the workshop.

TOPIC 1: *Introduction to Computer Forensics***.**

Q-1: Estimate the fraction of disk blocks affected by formatting a hard disk. 1) 100% 2) 65% 3) 20% 4) Less than 5%
Answer: *Less than 5%*

TOPIC 2: *File Systems***.**

Q-2: In order to access a file from a FAT filesystem, what information is absolutely necessary? 1) Name and ending address of file content 2) Name, file size, and ending address of file content 3) Name and starting address of file content 4) File size and starting address of file content 5) None of the above
Answer: *Name and starting address of file content*

Q-3: If sector 0 is lost/damaged in FAT12/16, what problem does it cause? 1) The volume's sector, cluster, etc. sizes cannot be determined 2) The volume's maximum file sizes become unavailable 3) The number of file allocation tables becomes unknown 4) More than one of the above 5) None of the above
Answer: *More than one of the above ("The volume's sector, cluster, etc. sizes cannot be determined" and "The number of file allocation tables becomes unknown")*

TOPIC 3: *File Carving***.**

Q-4: File carving is especially useful in which of the following one or more situations? 1) An operating system drive is examined as an external drive 2) Many potentially desired files have been deleted 3) Files have recently been defragmented 4) The file allocation table on a FAT file system has been corrupted 5) More than one of the above
Answer: *More than one of the above ("Many potentially desired files have been deleted" and "The file allocation table on a FAT file system has been corrupted")*

Q-5: File carving is the most effective in which one or more of the following scenarios? 1) Drive is highly fragmented 2) Drive is recently defragmented 3) System used to examine drive has low space 4) System used to examine drive has high space 5) More than one of the above
Answer: *System used to examine drive has high space*

TOPIC 4: *Windows Registry.*

Q-6: Which of the following actions is best described as an example of registry forensics? 1) An investigator uses Volatility to examine a file 2) An investigator uses LiME to access a memory dump 3) An investigator reviews the SAM hive to obtain password hashes 4) An investigator examines access timelines using FTK
Answer: *An investigator reviews the SAM hive to obtain password hashes*

Q-7: A USB drive with an unknown owner is found in a corporate setting. How might a forensic investigator typically determine whether that particular drive was plugged into any given Windows machine? 1) Examine all ntuser.dat files to determine if a user plugged it into the machine 2) Check the system registry file to see if it was plugged into that machine 3) Check the software registry file to see if it has been used by any particular piece of software 4) More than one of the above 5) None of the above
Answer: *Check the system registry file to see if it was plugged into that machine*

8 Evaluation

8.1 Data Collection Instruments

We use three instruments to measure students' learning gain and their interest in peer instruction: 1) quiz and 2) the clicker responses of peer instruction questions, and 3) survey.

Quiz. Three quizzes are used; the first covers the first two topics and the other two cover the later two topics. Each quiz is presented twice to students–once before and once after a topic is covered (refer to Figure 5). The ease of the questions is intended, as these questions are not conceptual, but they are used rather to determine how students retained knowledge from each of the lecture sections. The full content of the quizzes is included in Appendix A.

Clicker Response. We also collect data on student responses via clickers on peer instruction questions before and after discussion of each question (refer to Figure 4). It is useful to measure immediate impact on student learning.

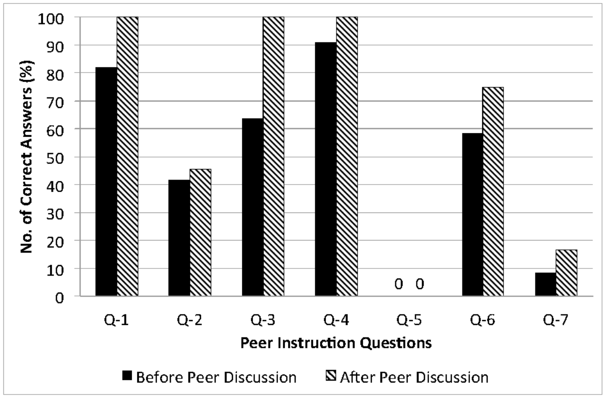

Figure 4: Percentage of correct answers of peer instruction questions (Q) before and after the peer discussions

Survey. We also provide the results of a student attitudinal survey of workshop participants (refer to Figure 3). The survey instrument was provided by Beth Simon and Leo Porter of UC San Diego, and Cynthia Lee of Stanford University. Results from this survey instrument have been published for numerous peer instruction courses, some of which have been extensively studied for effectivness by a variety of measures, providing useful comparisons for our initial workshop-based pilot of these materials (*e.g.*, [5][8][9]).

The peer instruction survey gathered information on prior usage of clickers, workshop preparation (reading material and quiz), peer discussion, clicker usage, and lecture pacing. It uses a number of questions with a Likert scale to determine lecture preparation as well as opinions on peer instruction, clicker usage, and attentiveness; there are also a few additional opinion questions discussing details specific to the iteration of the workshop (refer to Table 4 in Appendix).

8.2 Evaluation Results

Clicker Response. Figure 4 shows the percentage of correct answers to peer instruction questions. It presents a comparison of each question when it is presented to students twice–before and after the peer discussion. The comparison shows visible improvement in correct answers across the entire set of question pairs except question 5.

We assess question 5 carefully and try to determine why none of the students could reply correctly. We believe that it may be due to the wording of the question or its answers.

We also obtain some feedback from the students during the workshop that suggested some attendees may have felt that the question was a comparison of file carving vs. other unnamed evidence collection techniques rather than a comparison of potentially ideal usage sit-

Table 3: Peer instruction lecture preparation, peer instruction, and clicker usage opinions

Question	Average Opinion
Thinking about clicker questions on my own, before discussing with people around me, helped me learn the workshop material.	87%
I read the required material before the workshop.	89%
The pre-workshop reading quiz helped me recognize what was difficult in the reading.	76%
Most of the time my group actually discussed the clicker question.	87%
Discussing course topics with my seatmates in the workshop helped me better understand the workshop material.	96%
The immediate feedback from clickers helped me focus on weaknesses in my understanding of the workshop material.	91%
Knowing the right answer is the only important part of the clicker question.	49%
Generally, by the time we finished with a question and discussion, I felt pretty clear about it.	80%
Clickers are an easy-to-use class collaboration tool.	89%
Clickers helped me pay attention in this workshop compared to traditional lectures.	82%
Using clickers with discussion is valuable for my learning.	80%
I recommend that other instructors use this approach (reading quizzes, clickers, in-class discussion) in their courses.	91%

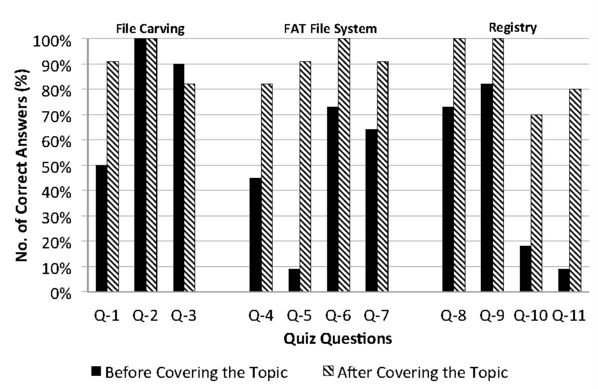

Figure 5: Percentage of correct answers of quiz questions (Q) before and after a topic is covered.

uations for file carving. Also, there was confusion as to whether the multiple-choices presented to students referred to the situation on an investigator's or a suspect's drive as well as "space" meaning "free space" in the case of answers c and d.

Quiz. Figure 5 summarizes the results of three quizzes and presents a comparison of each quiz question taken by the students before and after the respective topic is covered. It shows the evidence of student learning gain since it shows clear improvement across all the questions.

Survey. Table 3 presents the results of the student attitudinal survey portion of the workshop evaluation. It shows that the most of the students find it useful to think about a clicker question before discussing it with other students and the discussion helps them understand the concept better. 91% of students would recommend peer instruction be adopted by other instructors, a number that corroborates results reported for this same survey instrument elsewhere in the computer science education literature, for example: 91% average across 7 CS1 (introductory programming) courses [8], 88% average across 6 comptuer architecture and theory of computation courses [5], and an average of 91% in 10 computer science courses at small liberal arts colleges [9]. This close corroboration is representative of the results for the other questions in the survey.

Table 4 in Appendix summarizes the students opinion about the peer instruction workshop. It shows that the students have a generally positive experience of the workshop. They have adequate time to understand the questions and vote for the correct answer. 89% students agree that the allowable duration for group discussions is sufficient.

9 Conclusion

We developed 108 peer instruction questions for a digital forensics curriculum. The questions utilized concept triggers including "compare and contrast", "deliberate ambiguity" and "trolling for misconceptions". They used four question presentation types i.e. scenario, example, feature, and definitional.

We held a a four-hour long workshop on computer forensics to test the peer instruction methodology and a subset of peer instruction questions. The workshop, while small, shows that peer instruction holds promise in its implementation in digital forensics courses. The participant students provided positive responses in the peer instruction and clicker survey. They showed visible learning gains in nearly all of the quiz and peer instruction questions used. Our average workshop learning gains for quiz and peer instruction questions are 34% and 13% respectively.

References

[1] Beatty, I. D., Gerace, W. J., Leonard, W. J., and Dufresne, R. J. Designing effective questions for classroom response system teaching. *American Journal of Physics 74*, 1 (2006), 31–39.

[2] Crouch, C. H., and Mazur, E. Peer Instruction: Ten years of experience and results. *American Journal of Physics 69* (Sept. 2001), 970–977.

[3] Esper, S. A discussion on adopting peer instruction in a course focused on risk management. *J. Comput. Sci. Coll. 29*, 4 (Apr. 2014), 175–182.

[4] Johnson, W. E., Luzader, A., Ahmed, I., Roussev, V., III, G. G. R., and Lee, C. B. Development of peer instruction questions for cybersecurity education. In *2016 USENIX Workshop on Advances in Security Education (ASE 16)* (Austin, TX, 2016), USENIX Association.

[5] Lee, C. B., Garcia, S., and Porter, L. Can peer instruction be effective in upper-division computer science courses? *Trans. Comput. Educ. 13*, 3 (Aug. 2013), 12:1–12:22.

[6] Liao, S. N., Zingaro, D., Laurenzano, M. A., Griswold, W. G., and Porter, L. Lightweight, early identification of at-risk cs1 students. In *Proceedings of the 2016 ACM Conference on International Computing Education Research* (New York, NY, USA, 2016), ICER '16, ACM, pp. 123–131.

[7] Porter, L., Bailey Lee, C., and Simon, B. Halving fail rates using peer instruction: A study of four computer science courses. In *Proceeding of the 44th ACM Technical Symposium on Computer Science Education* (New York, NY, USA, 2013), SIGCSE '13, ACM, pp. 177–182.

[8] Porter, L., Bouvier, D., Cutts, Q., Grissom, S., Lee, C., McCartney, R., Zingaro, D., and Simon, B. A multi-institutional study of peer instruction in introductory computing. In *Proceedings of the 47th ACM Technical Symposium on Computing Science Education* (New York, NY, USA, 2016), SIGCSE '16, ACM, pp. 358–363.

[9] Porter, L., Garcia, S., Glick, J., Matusiewicz, A., and Taylor, C. Peer instruction in computer science at small liberal arts colleges. In *Proceedings of the 18th ACM Conference on Innovation and Technology in Computer Science Education* (New York, NY, USA, 2013), ITiCSE '13, ACM, pp. 129–134.

[10] Porter, L., and Simon, B. Retaining nearly one-third more majors with a trio of instructional best practices in cs1. In *Proceeding of the 44th ACM Technical Symposium on Computer Science Education* (New York, NY, USA, 2013), SIGCSE '13, ACM, pp. 165–170.

[11] Roussev, V. *Digital Forensic Science: Issues, Methods, and Challenges.* Synthesis Lectures on Information Security, Privacy, & Trust. Morgan & Claypool Publishers, 2016.

Appendices

A Workshop Quizzes

File Systems Quiz.

Q-1: What are the two primary data structures of FAT systems?
1) Directory entry, File allocation table
2) Cluster entry, File allocation table
3) File entry, Quick allocation table
Answer: *Directory entry, File allocation table*

Q-2: How is the filename “conf.ini” stored in a FAT file system that utilizes short filenames?
1) `conf.ini`
2) `CONF.INI`
3) `CONFINI`
4) `CONF INI`
Answer: *`CONF INI`*

Q-3: FAT32 maintains a backup BIOS Parameter Block.
1) True
2) False
Answer: *True*

Q-4: How does a FAT file system denote the end of a file?
1) The number of the final cluster equaling the stored number of clusters (-1 for zero indexing)
2) A FAT entry marked EOF
3) The final FAT entry for the file is marked ”NULL”
4) Each file is allocated the same initial space, and the first ”NULL” entry in the file's allocation table is the first cluster following the end of file
Answer: *A FAT entry marked EOF*

File Carving Quiz.

Q-5: Traditional carving uses these to find potential files:
1) Known headers
2) Known filenames
3) Allocated clusters following blocks of unallocated clusters
4) Recovered file system metadata
Answer: *Known headers*

Q-6: Which of the following is a known issue with carving?
1) Fragmentation
2) Milestones
3) Unprintable bytes in headers
Answer: *Fragmentation*

Q-7: File carving could efficiently utilize distributed systems.
1) True
2) False
Answer: *True*

Windows Registry Quiz.

Q-8: What are the primary registry files known as?
1) Hives
2) Keys
3) Values
4) Root files
Answer: *Hives*

Q-9: What is the timestamp given to any registry key?
1) LastWriteTime
2) LastReadTime
3) CreatedTime
Answer: *LastWriteTime*

Q-10: Which of the following stores data in the registry?
1) Key
2) Value
3) Data
4) Hive
Answer: *Data*

Q-11: Where can user password hashes be found?
1) SYSTEM
2) SECURITY
3) SOFTWARE
4) SAM
5) DEFAULT
Answer: *SAM*

B Quiz for Pre-class Reading Material

1. Which of these is not a Windows registry root key? 1) HKEY_LOCAL_MACHINE
   2) HKEY_USERS
   3) HKEY_CURRENT_CONFIG
   4) HKEY_ROOT_USER
   5) HKEY_CLASSES_ROOT
   Answer: HKEY_ROOT_USER
2. Which of the following utilizes a bitmap to denote cluster allocation? 1) FAT16 2) NTFS 3) FAT32
   Answer: NTFS
3. Which of these carvers was built as a direct improvement to Foremost? 1) Photorec 2) Scalpel 3) Binwalk 4) FTK carver 5) Magic Rescue
   Answer: Scalpel
4. Which of these features of the registry is most known for using ROT13 for encoding data? 1) Autorun 2) MRU 3) UserAssist 4) USBSTOR
   Answer: UserAssist
5. Which web browser utilizes the TypedURLs registry key? 1) Mozilla Firefox 2) Internet Explorer 3) Google Chrome 4) Opera Browser
   Answer: Internet Explorer

C Survey on Workshop Opinion

Table 4: Workshop-specific opinions

<table>
<tr><td colspan="60">From the point of helping me learn, the content of clicker questions was</td></tr>
<tr><td colspan="12">Much too hard
0%</td><td colspan="12">Too hard
0%</td><td colspan="12">OK
100%</td><td colspan="12">Too easy
0%</td><td colspan="12">Much too easy
0%</td></tr>
<tr><td colspan="60">In general, the instructor gave us enough time to read and understand the questions before the first vote.</td></tr>
<tr><td colspan="12">No, far too little time
0%</td><td colspan="12">No, too little time
0%</td><td colspan="12">OK amount of time
89%</td><td colspan="12">Yes, too much time
11%</td><td colspan="12">Yes, far too much time
0%</td></tr>
<tr><td colspan="60">Which of the following best describes your discussion practices in this group?</td></tr>
<tr><td colspan="12">I always discuss with the group around me, it helps me learn
78%</td><td colspan="12">I always discuss with the group around me, I don’t really learn, but I stay awake
0%</td><td colspan="12">I sometimes discuss, it depends
22%</td><td colspan="12">I rarely discuss, I don’t think I get a lot out of it
0%</td><td colspan="12">I rarely discuss, I’m too shy
0%</td></tr>
<tr><td colspan="60">The amount of time generally allowed for peer discussion was</td></tr>
<tr><td colspan="12">Much too short
0%</td><td colspan="12">Too short
11%</td><td colspan="12">About right
89%</td><td colspan="12">Too long
0%</td><td colspan="12">Much too long
0%</td></tr>
<tr><td colspan="60">In general, the time allowed for class-wide discussion (after the group vote) was</td></tr>
<tr><td colspan="12">Much too short
0%</td><td colspan="12">Too short
11%</td><td colspan="12">About right
89%</td><td colspan="12">Too long
0%</td><td colspan="12">Much too long
0%</td></tr>
<tr><td colspan="60">In general, it was helpful for the instructor to begin class-wide discussion by having students give an explanation.</td></tr>
<tr><td colspan="20">N/A - The instructor rarely did this
11%</td><td colspan="20">It’s not helpful to hear other students’ explanations
0%</td><td colspan="20">It was helpful to hear other students’ explanations
89%</td></tr>
<tr><td colspan="60">The professor explained the value of using clickers in this class.</td></tr>
<tr><td colspan="15">Not at all
0%</td><td colspan="15">Somewhat, but I was still unclear why we were doing it
11%</td><td colspan="15">Yes, they explained it well
67%</td><td colspan="15">Yes, they explained it too much
22%</td></tr>
</table>