CS4677, Digital Forensics (3-2) -- Syllabus for Winter 2026

Catalog description

This course covers the fundamentals of computer forensics in the context of DoN/DoD information operations. Students examine how information is stored and how it may be deliberately hidden and/or subverted. Coverage includes: practical forensic examination and analysis, techniques of evidence recovery, legal preparation of evidence, common forensic tools, principles of original integrity, disk examination, and logging. Prerequisite: CS3600.

Instructor

Prof. Neil Rowe, ncrowe@nps.edu, (831) 656-2462 (office), GE-328, https://faculty.nps.edu/ncrowe. Use that email address for submissions of the homework and project; do not mail to the Sakai instructor account. No official office hours, so send him an email if you want to meet in person.

Classes

Lectures are Monday, Tuesday, Wednesday, and Thursday at 1000. We will only use one lab hour per week, as the hour lab hour is for your class project and done on your own; we will only use the extra lab hour on Friday for student presentations in the second-to-last week. Like most classes at NPS, we expect you to spend at least six hours per week outside of class, for a commitment of 10 hours per week. We will record a video of the classes which will be posted for around a month in Sakai.

Grading

Grades are based on three homework assignments, three quizzes, and a class project; the quizzes are weighted less than the other grades. Guidelines for the class project are given later in this syllabus. The median course grade is an A-. Scores on homework and quizzes average around 70%, so these are graduate-school questions, not training-course questions.

Homework should be submitted by email to ncrowe@nps.edu as either a Microsoft Word docx file (which is preferred) or an Acrobat PDF file if you do not use Word. It should be submitted as a single document with your name included in the file name, your name at the top of the document, all material for each problem together (no appendices), and without the text of the problems (just your answers). Homework questions are written so quoting or paraphrasing something written elsewhere will not answer the question. Homework must be done by each student on their own without consulting anyone besides the instructor. Large language models like ChatGPT may not be used to produce homework answers. The penalty is 15% for assignments received after the due date, but homework is not accepted after solutions are made available.

Quizzes will be given in class on the dates listed on the schedule. They are 50 minutes long, and are open-book and open-notes, meaning you can refer to paper copies of the book, the slides, and your notes during the quiz. You may not use any electronic devices during the quiz.

Guidelines for the project are given at the end of this syllabus.

Textbook and class materials

We use a site in Sakai (cle.nps.edu); tell the instructor if you cannot log in to it. The instructor is writing a textbook, and a draft is posted on the Sakai site under the "Resources" tab. You should read it as listed below, except you can skip the exercises when not assigned. We also post under "Resources" the PowerPoint slides shown in class, the homework, and this syllabus. The initial full set of slides is also posted in a single zip-format file. However, there will be small updates to the slides during the course that will be posted in Sakai.

Schedule

We will also prepare a schedule for student presentations in March.
By 1/9: Read Chapter 1 in the textbook
By 1/16: Read Chapter 2
1/19 (Monday): Holiday
By 1/23: Read Chapter 3
On 1/26 (Monday): Homework 1 due online by class time
By 1/29: Read Chapter 4
On 1/29 (Thursday): Quiz 1 in class on chapters 1-4 of the book
On 2/3: Propose orally in class your class-project topic
By 2/6: Read Chapters 5 and 6
By 2/13: Read Chapters 7 and 8
2/16 (Monday): Holiday
On 2/23 (Monday): Homework 2 due online by class time
On 2/26 (Thursday): Quiz 2 online in Sakai, on chapters 5-8 of the book
By 2/27: Read Chapter 9
By 3/6: Read Chapters 10 and 11
By 3/11: Read Chapter 12
On 3/11 (Wednesday) student project presentations
On 3/12 (Thursday): Homework 3 due online by class time; student project presentations
On 3/13 (Friday) student project presentations, 0900-0950
By 3/16 (Monday): Read Chapter 14 (we skip Chapter 13);
On 3/18 (Wednesday): Quiz 3 in class on chapters 9-12 and 14 of the book; last day of class
On 3/20 (Friday): Project writeup due by 1200 PST to ncrowe@nps.edu

Topics covered in order in the course:

- Introduction: Application areas, cybercrime types, e-discovery, differences from traditional forensics, types of forensic media (computers, devices, and servers), Locard's Principle, other forensic principles, forensic tools
- Forensic investigation procedures: Digital crime schemes, recording of findings, the role of inference
- Basic concepts needed from computer science: Time required for basic iterations, text encoding, number and string encoding, data compression, tables and trees, data hashing and signatures, Bloom filters, entropy cosine similarity, F-scores
- Forensic data acquisition: Different sources of forensic data, method options, hardware for acquisition, live-system acquisition, volatile memory issues, secondary-storage acquisition, mobile device acquisition, copying methods, forensic containers, write blockers, virtual machines, Web crawlers, network packet collection
- Forensic duplication: Types of drive images, full imaging versus sparse imaging, dealing with blank space, unallocated and slack space, hidden files, secondary-storage terms, line terminators,

endianness, original copies, chain of custody, evidence tags, anonymization, secure erasure, conditions for long-term media storage, password and key recovery, brute-force and dictionary attacks, rainbow tables, hardware manipulation

- File systems: Drive volumes, file metadata, evidence aggregation, timestamp analysis, classification of users from metadata, finding anomalous files, host-based intrusion detection, Bayesian methods for searching for malware
- File types: The SleuthKit tool, retrieving files from a drive image, typical file-type distributions, hex editors, text files, encoded data, email and message files, determining data length, Naïve Bayes classification of file fragments, classification of fragments by entropy and cosine similarity, matching byte sequences between files, known file filters like NSRL, extensions to NSRL from corpus analysis
- Text analysis: String search algorithms, comparing byte sequences, comparing word and character histograms to find similar data, establishing statistical significance of classification, stopwords, destemming, part-of-speech taggers, word-sense classification, determining the language of file names and contents
- Investigative processes: Autopsy, malware investigation, rootkit installation investigation, hacking investigation, intrusions, espionage, fraud, cyber stalking, identity theft, root-cause analysis
- File carving and personal artifacts: Recognizing and connecting fragments, finding encoding schemes, rating email-address candidates found, finding personal names, finding other useful forensic artifacts
- Registry and times: Contents of the Windows Registry, registry-analysis tools, timestamp updating, timeline analysis, evolution of software versions
- Networking forensics: Following leads across machines, server and cloud forensics, Web site forensics, measuring relationships between machines, measuring relationships between people, visualizing social networks
- Anti-forensics: Types of adversaries, concealment, obfuscation, system and hardware manipulation, static clues, suspicious software, deception in media, general deception theory
- Legal aspects of forensics: Cybercrime laws, privacy laws, intellectual property laws, precedents in case law, search authorization for law enforcement, authentication of digital evidence, affidavits, court testimony with digital forensics, digital forensics in civil cases, cyberwarfare forensics

Guidelines for the class project

General guidelines:

- The instructor must approve the topic. You will describe your proposed topic in class at a session a few weeks into the course. You may also be asked to give a progress report later.
- It should involve at least 30 hours of effort by each student.
- It can be a group project with up to three people. Then you must identify in the writeup who contributed to each part to the project. Some individual work by each member of the group must be identifiable.
- It can involve programming, data analysis, a survey of other people's ideas, or some combination of the three. But it must include either some programming or some data analysis.
- It should focus on material related to the digital forensics topics covered in the course. Since we are not going to cover much about networking, a networking topic would only be acceptable if it involves looking at the data or packets transmitted over a network.

- It should not be exactly the topic of your thesis or dissertation if you have one, but can be something related to the topic. For instance, it may be something listed as "future work" in the thesis.
- You give a presentation in class on your project towards the end of the course. The work may not be complete at the time of the presentation, but try to explain what you are trying to do.
- The project will be graded separately on the presentation, writeup, testing, and ambition. Some testing of the ideas you report is required. The quality of your English is not graded in the presentation or writeup as long as you are reasonably understandable.

For the writeup:

- Submit the project writeup electronically to ncrowe@nps.edu.
- It should be at least 3000 words for one student, 4000 words for two students, and 5000 words for three students. Give the full names of the students on the first page.
- No particular format is required, but be sure to cover background, previous work if applicable, what you did, what results you obtained, and conclusions about your project.
- Wording must be your own, except for occasional quotations for which you should use quotation marks. Large language models such as ChatGPT may not be used to generate text.
- If it is a group project, identify what each person contributed to the project.
- If the project relates to your thesis or another project you have done, explain how and what is different.
- Give some background on the technique or problem addressed and explain why it is important.
- Give some good-quality references to related previous work by other people (not just Web sites). Try to offer insights rather than just repeating what other people said about the previous work.
- If you used someone’s data, explain where you got it, like the name of its Web site. If you created your own data, say so.
- For the program or data analysis, submit results from it and show any code or queries you used to obtain it. Graphs and tables are important to include to help summarize data.
- Discuss future directions students could take in working on your subject.
- Submit it by email as a file in Microsoft Word or PDF format.

Some ideas for the final project:

- Analyze a subset of drives to determine their usage. It is easiest to purchase drives online (Ebay has some), though bear in mind some will be erased or faulty. There are also some drives available for free from NIST.
- Analyze how things you do on a computer or phone indirectly affects its files.
- Analyze a particular file type or set of types in our corpus to determine its usage and study its data. We have tools to compare files.
- Analyze a cell phone of your own using an open-source tool. However, your rate of success can be highly variable depending on the tool. An alternative is to analyze the backup storage for phones. Some vendors have tools for their phones, and there are a few general tools like Android Device Bridge and AFLogical OSE. If you are willing to “jailbreak” a phone, you can get more access to it, but try this on an old phone you are not using much.
- Image NPS classroom and laboratory machines, and extract new kinds of data.
- Analyze the occurrence of malware in the corpus.
- Analyze the full set of effects on a drive of particular malware attacks.
- Do forensics on drones or vehicles.
- Do forensics on cyber-physical systems like power plants and building-management systems.

- Test forensic methods for analyzing virtual machines.
- Analyze patterns of timestamps in particular directories across drives; we can supply data from our corpus.
- Collect data from Web pages using tools like the Python Scrapy module and HTML parsers.
- Analyze cookies left on drives.
- Analyze browser data; possible tools are Web Historian, Index.Data Analyzer, Chromensics, Image Cache Viewer, and Browsing History Viewer.
- Analyze the keyword search terms used on particular drives, and try to infer interests of the users.
- Analyze Dark Web anonymous data.
- Analyze statistics of malware files. We do have the malware from our corpus, but it doesn't have much variety; it might be better to get examples elsewhere.
- Develop statistical methods to find the boundaries of the major partitions of a file.
- Develop good methods for comparing drives to find similar ones, by studying subsets of drives and of clues to match.
- Trace the history of updates to a particular software file. This is particularly useful for executable files to detect malware by its unusual timestamps and dissimilarity to legitimate files.
- Build social networks from drives that have similarities, then use graph algorithms to characterize central nodes and key paths.
- Develop new techniques for visualizing the data in our corpus, especially social networks.
- Develop big-data techniques like Hadoop for handling a large forensic corpus.
- Study effectiveness of a memory acquisition tool like Volatility.
- Analyze particular methods for impeding forensics on a drive ("antiforensics").
- Test our identified uninteresting files to see how often our analysis is correct.
- Classify new file extensions and directory names that have happened in the last five years.
- Develop methods to detect leakage of classified information by looking for keywords.
- Determine the dialect of text.

Look at the "Sources of Forensic Data" paper on the Sakai site for ideas about where to get data.