CSCI E-11

Artificial Intelligence, the Internet-of-Things, and Cybersecurity

Spring 2026

Class 1

January 27, 2026

Harvard
Extension School
HARVARD DIVISION OF
CONTINUING EDUCATION

Welcome to CSCI E-11: AI, IoT and Cybersecurity!

Simson Garfinkel
Visiting Lecturer
Please contact through Canvas

Geoff Cohen
Senior Technologist, PCLOB
geoff_cohen@fas.harvard.edu
(But please use Canvas!)

Lecture slides with notes available as PDFs from the Canvas Files tab.

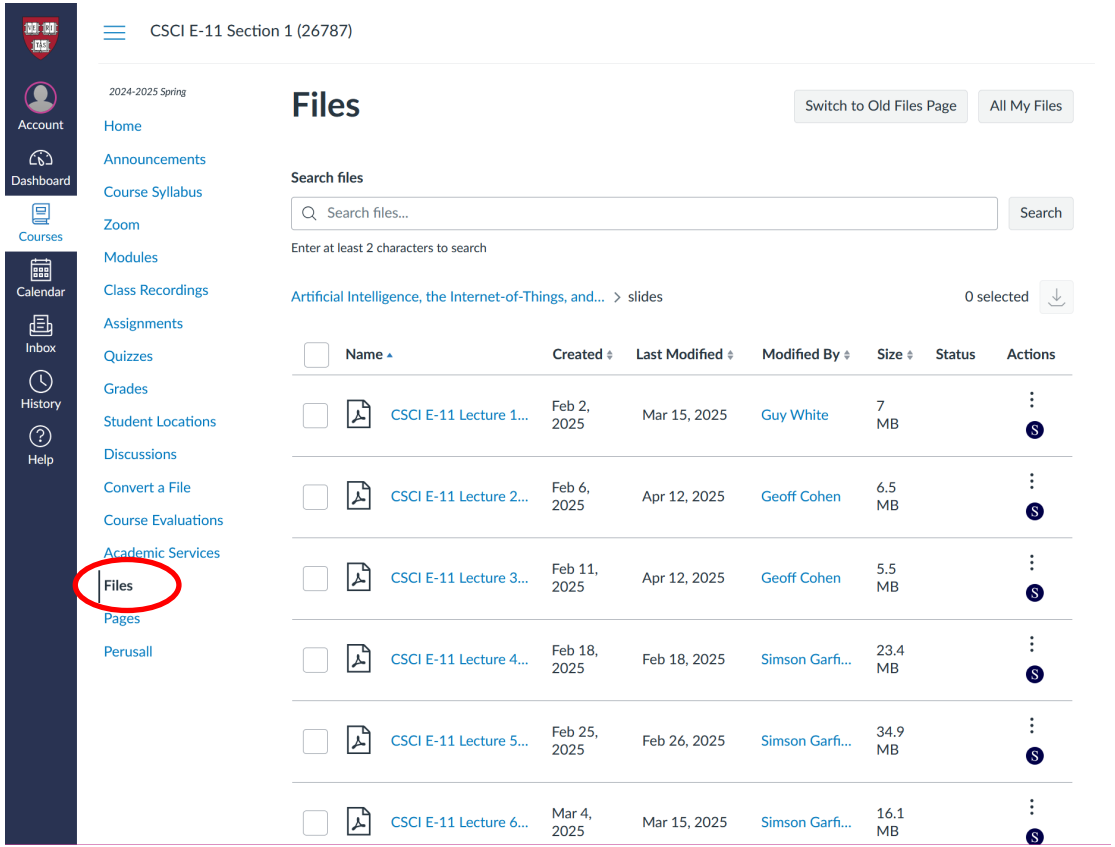

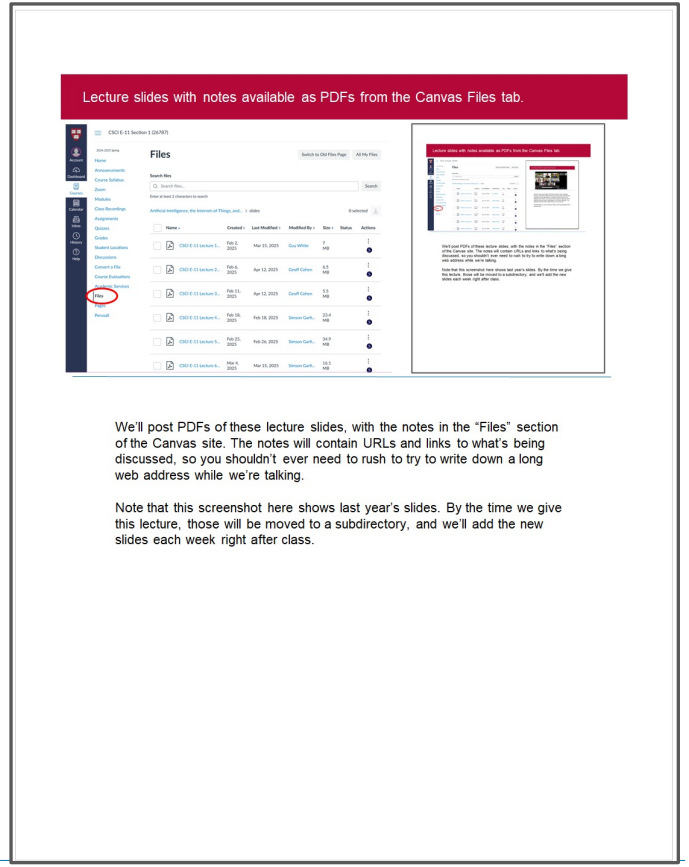

This is an Immersive Class!
Use the Zoom Chat+ while watching Immersive Classroom.

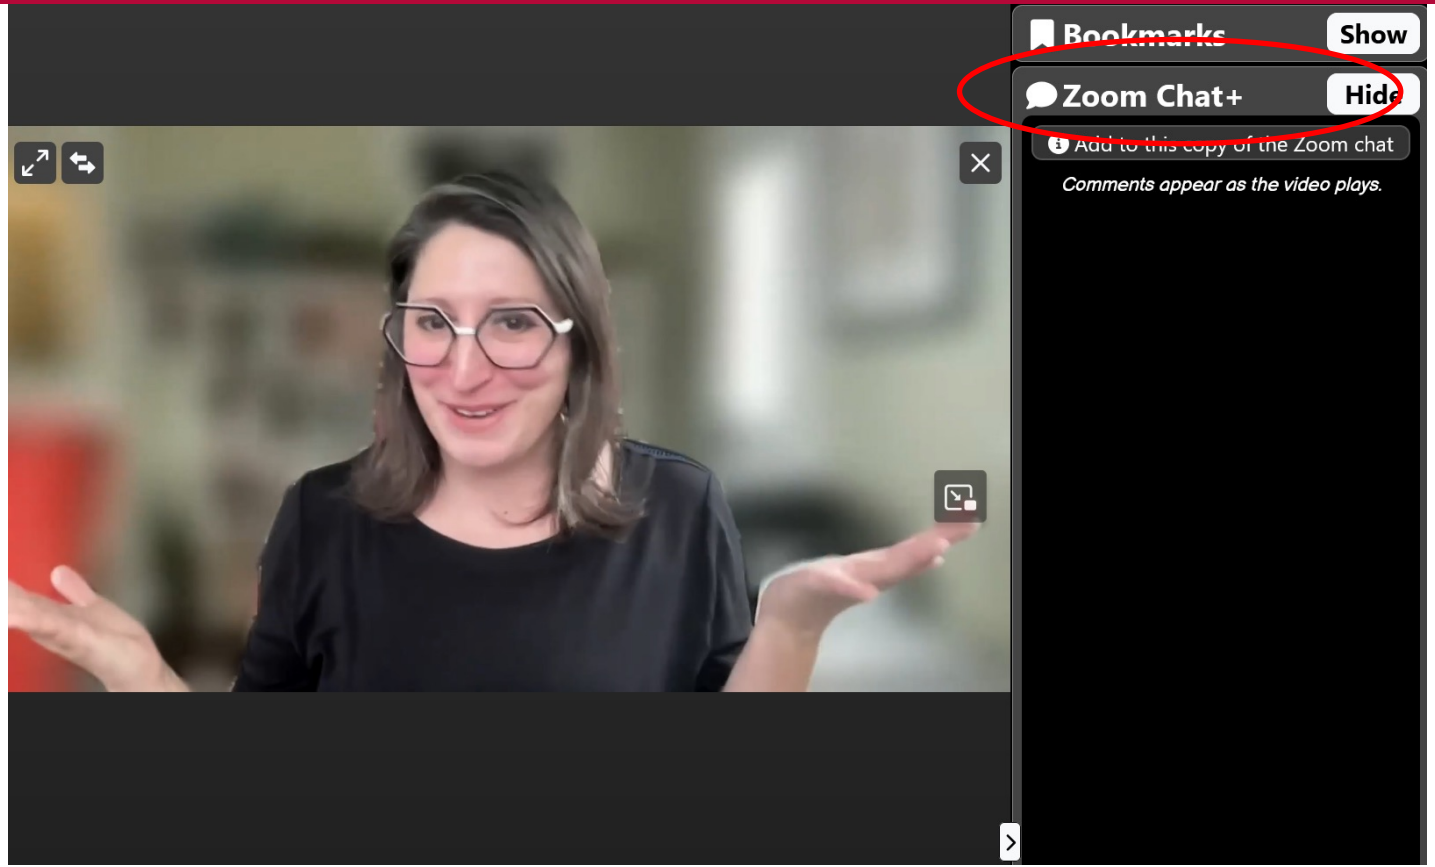

You can also chat live with other Immersive Classroom watchers

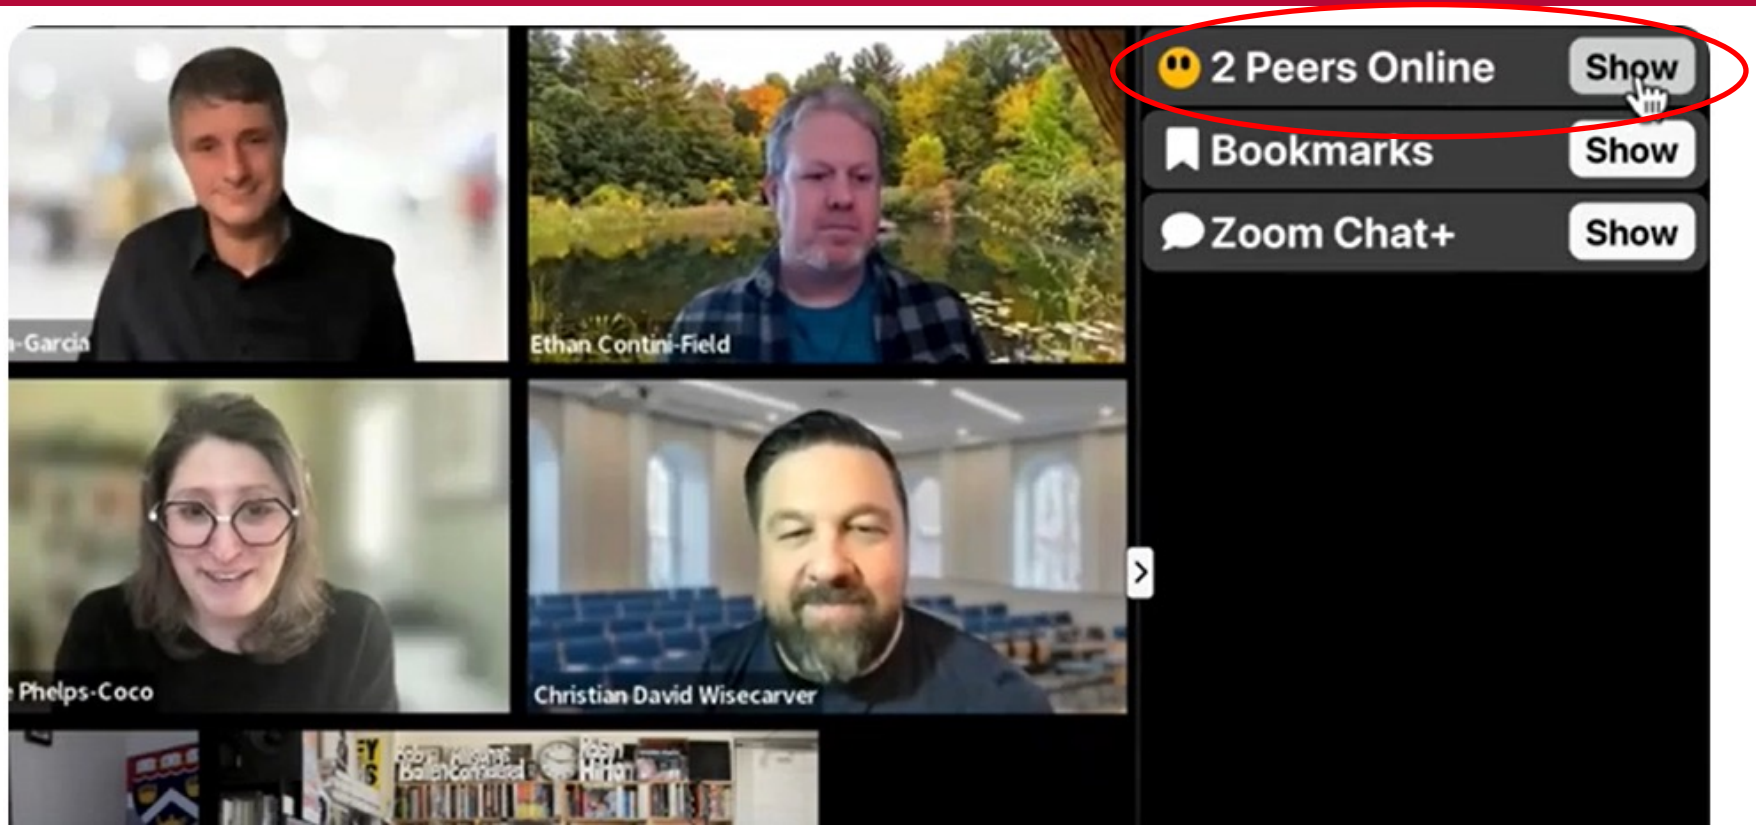

Let's try it out! What was your first computer? What was your earliest computer experience?

Simson Garfinkel with his first computer (1978)

Geoff with his first computer (1983-84?)

Let’s start with a case study of a security issue in an IoT device.

In 2017, Amazon introduced “Cloud Cam,” an Internet camera with Alexa.

In 2022, Amazon announced it would no longer support the Cloud Cam.

To our valued Cloud Cam customer,

With your help over the last five years, Cloud Cam has served as a reliable indoor security camera and a hub for Amazon Key-compatible smart locks that work with Alexa. As the number of Alexa smart home devices continues to grow, we are focusing efforts on Ring, Blink, and other technologies that make your home smarter and simplify your everyday routines. Therefore, we have decided to no longer continue support for Amazon Cloud Cam and its companion apps.

What this means for you:
On December 2, 2022, you will no longer be able to use your Cloud Cam device or its companion apps. Until then, you will be able to download any video recordings if available. All video history will be deleted on December 2, 2022. Also, if you are using Cloud Cam – Key Edition as a Zigbee hub, you will no longer be able to use it to connect to compatible locks or manage pin codes in the Amazon Key app.

In 2025, Amazon disclosed a vulnerability that allowed attackers to gain control.

CNA: Amazon

Published: 2025-06-12 **Updated:** 2025-10-14

Title: Insecure device pairing in end of life Amazon Cloud Cam

Tags: unsupported-when-assigned

Description

Amazon Cloud Cam is a home security camera that was deprecated on December 2, 2022, is end of life, and is no longer actively supported. When a user powers on the Amazon Cloud Cam, the device attempts to connect to a remote service infrastructure that has been deprecated due to end-of-life status. The device defaults to a pairing status in which an arbitrary user can bypass SSL pinning to associate the device to an arbitrary network, allowing for network traffic interception and modification. We recommend customers discontinue usage of any remaining Amazon Cloud Cams.

CWE 1 Total

Learn more

- CWE-672: CWE-672 Operation on a Resource after Expiration or Release

AI capabilities increase the potential impact of the vulnerability.

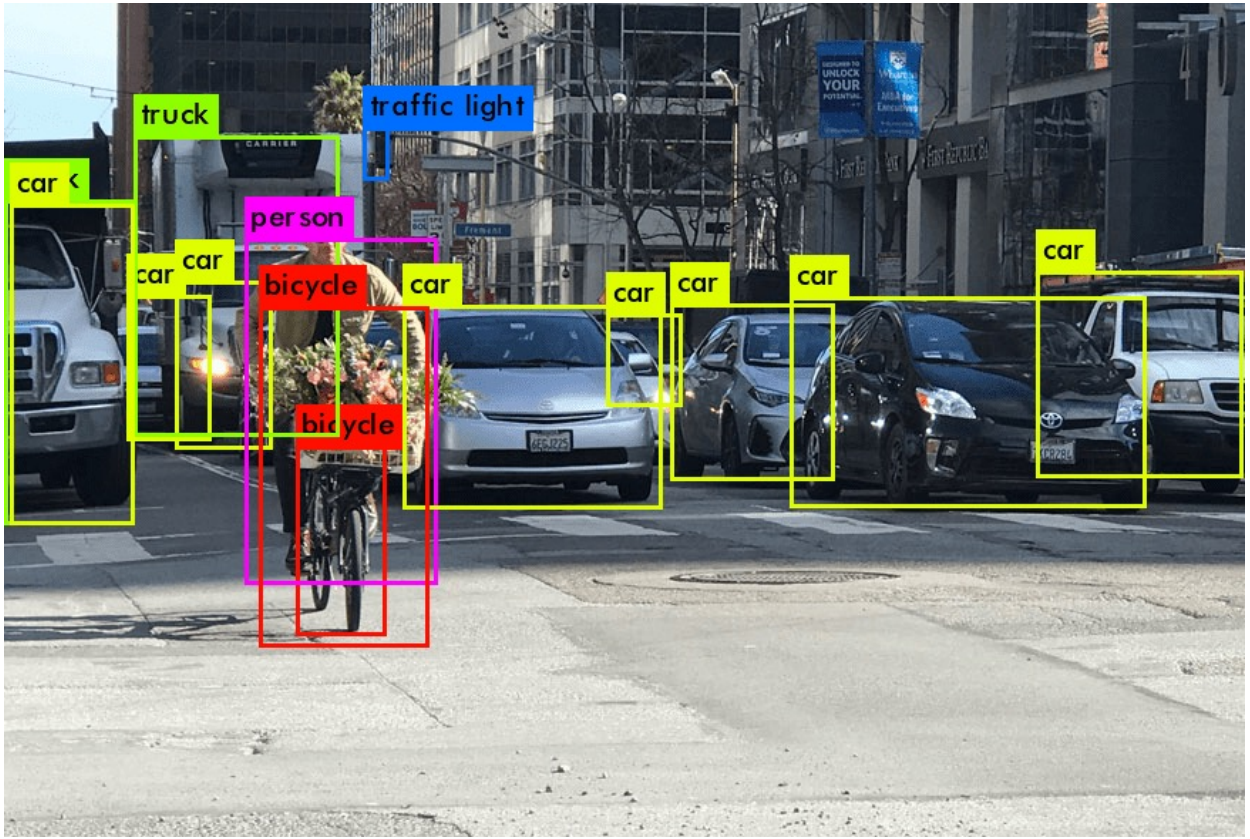

Please fill out the two surveys.

CSCI E-11 Welcome Survey

This anonymous survey will help us gauge the skills and interests of the students. The results will be displayed during the first class!

Which courses or subjects have you previously taken (or are familiar with the material)?

- [ ] Programming (e.g., CSCI E-3 Introduction to Web Programming using JavaScript, CSCI E-22 Data Structur…
- [ ] Systems programming (e.g., CSCI E-28 Unix/Linux Systems Programming, Networking, Operating System…
- [ ] Introduction to Computer Science (e.g., CSCI E-7 Introduction to Computer Science with Python, CSCI E-1 …
- [ ] User Experience (e.g., CSCI E-34 User Experience Engineering or equivalents)
- [ ] Cloud computing (e.g., CSCI E-49B Cloud Architectures, Security, and Governance, CSCI E-94 Fundament …
- [ ] Data science (e.g., CSCI E-106 Data Modeling, CSCI E-109B Advanced Topics in Data Science, CSCI E-265…
- [ ] Artificial Intelligence (e.g., CSCI E-278 Applied Quantitative Finance and Machine Learning)
- [ ] CSCI E-184 Data Science and Artificial Intelligence: Ethics, Governance, and Laws
- [ ] Digital Media (e.g., DGMD E-1 Digital Media: From Ideas to Designs and Prototypes, DGMD E-11 Digital M …

Are you enrolled in (or planning to enroll in) any degree or certificate programs at Harvard or another institution?

- ○ Undergraduate Degree in Computer Science
- ○ Graduate Degree in Computer Science
- ○ Certificate in Cybersecurity
- ○ Other:

CSCI E-11 Spring 2026 SA-6

This is the Security Assessment Six Question Assessment (SA-6).

Please answer each question truthfully.

Your score is the sum of all the answers.

We will record the anonymized class scores at the beginning and end of the semester and report the results in class.

Generally, I diligently follow a routine about security practices

	1	2	3	4	5	
Strongly Disagree	○	○	○	○	○	Strongly Agree

I always pay attention to experts' advice about the steps I need to take to keep my online data and accounts safe.

	1	2	3	4	5	
Strongly Disagree	○	○	○	○	○	Strongly Agree

I am extremely knowledgeable about all the steps needed to keep my online data and accounts safe.

	1	2	3	4	5	
Strongly Disagree	○	○	○	○	○	Strongly Agree

Who Are You? (postponed until next week!)

Results of diagnostic survey

Results of SA-6

Course Management 1

Harvard's Canvas is the central point of our class.

Canvas can be accessed directly at Canvas.Harvard.edu or by logging into MyDCE and clicking the link to Canvas from your Dashboard.

To log in to the class Canvas page, you must use your HarvardKey.

If you do not have a HarvardKey, you may claim one at Key.Harvard.edu.

Please note, your HarvardKey login credential is different from your DCEKey.

Getting Help – Email, Office Hours, or use the Discussion Board!

Email us or attend (virtual) office hours

Simson: TBA/by appt
Geoff: Thursdays 5:30 pm*
or by appt
* but not this week

TAs

Guy White
guyewhite@cs50.harvard.edu
Matthew Lena
matthew.lena@hotmail.com

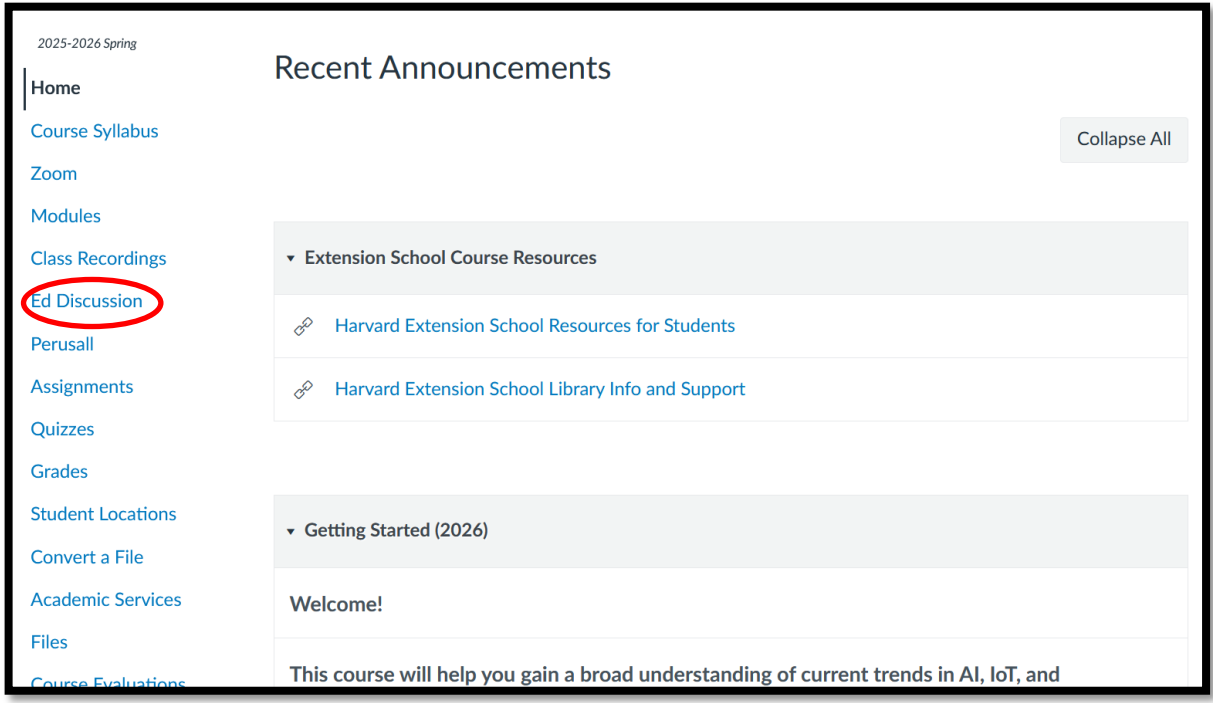

Discussion board

Who Are You, Part 2: Pre-Requisites (Official Version)

The official pre-requisites for this class are:

Any of these:

- CSCI E-3: Introduction to Web Programming Using JavaScript
- CSCI E-7: Introduction to Computer Science with Python
- CSCI E-10A: Introduction to Computer Science Using Java I

And:

- Familiarity with precalculus

Who Are You, Part 2: Pre-Requisites (Actual)

We hope that you are:

- Generally comfortable using the command line.
  - *—If you aren't, you will be!*
- Familiar with programming, especially Python
- Willing to read technical papers
- Willing to learn

Course Overview

Learning Goal #1: Subject matter knowledge

- The basic building blocks of engineering secure systems, e.g.: encryption, access control and privilege, identity and authentication.
- Various approaches to AI, and the foundations of neural net-based systems.
- The various architectural elements and constraints of the Internet of Things: power, sensors, networking, processing.

Learning Goal #2: Experience with risk management

- Learn about risks and threats to information systems
- Assessing risk likelihood and impact.
- How to design systems to minimize risks and to support the detection of, response to, and analysis of incidents.

Learning Goal #3: Practical experience

- Using Amazon Web Services, including starting and stopping virtual machines in the cloud, setting billing alerts, appropriately sizing VMs, and understanding VM and Amazon S3 security models.
- Setting up secure interactive server applications.
- Experience programming and configuring IoT devices.

What this class isn't

- Not a programming class (although you should be able to read and write Python).
- Not a software engineering class (although many of the things we talk about are important to good software engineering)
- Not a system administration class (although you'll learn some useful skills when administering your AWS instance)
- Not a public policy class (although Simson is teaching one this semester!)

This class will provide a high-level overview of topics in security, AI, and IoT.

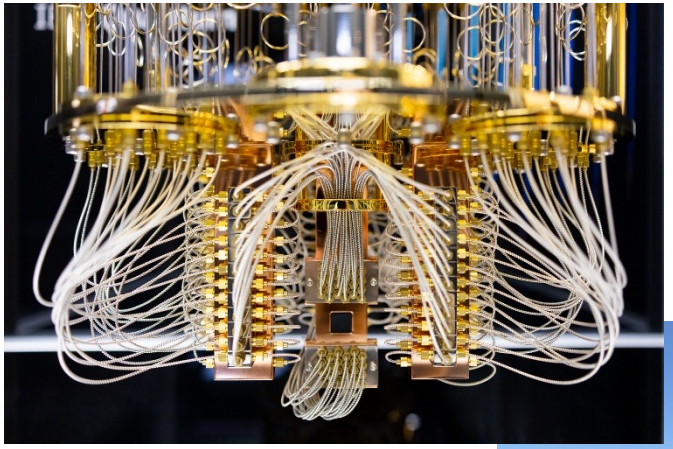

IBM Quantum Computer

Shake Alert Sensor

Big Dog Robot (2007)
Boston Dynamics

Modern (socio-) technical systems are extremely vulnerable to serious attacks.

The Hacker News

Cybercriminals Abuse Google Cloud Email Feature in Multi-Stage Phishing Campaign

Jan 02, 2026 Ravie Lakshmanan

The Hardest-Working Paper in America | Sunday, January 4, 2026

CHICAGO SUN-TIMES

Health care data breach affects 600,000 patients, Illinois agency says

Information such as addresses and case numbers were publicly viewable for around 32,000 customers with the Illinois Department of Human Services' division of rehabilitation services and another 670,000 Medicaid and Medicare Savings Program recipients, the agency said.

By Kade Heather | Jan 2, 2026, 4:38pm EST

The Record.
Recorded Future® News

University of Phoenix says 'numerous individuals' impacted by Oracle EBS breach

The news is full of incredible claims about AI's capabilities.

STAT+ | HEALTH TECH

FDA's new guidance on AI in drug development centers the risk introduced by the technology

The guidance is seen as the agency's attempt at regulatory clarity, a key barrier to greater AI use

AI means the end of internet search as we've known it

Despite fewer clicks, copyright fights, and sometimes iffy answers, AI could unlock new ways to summon all the world's knowledge

By Mat Honan

January 6, 2025

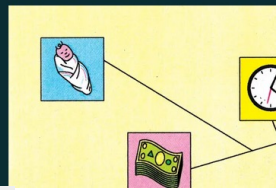

AI's next leap requires intimate access to your digital life

Tech executives are pitching "AI agents," digital helpers that can perform tasks on your computer, as the next big thing in artificial intelligence.

The Internet of Things promises to change industry, healthcare, and consumer products (and everything else)

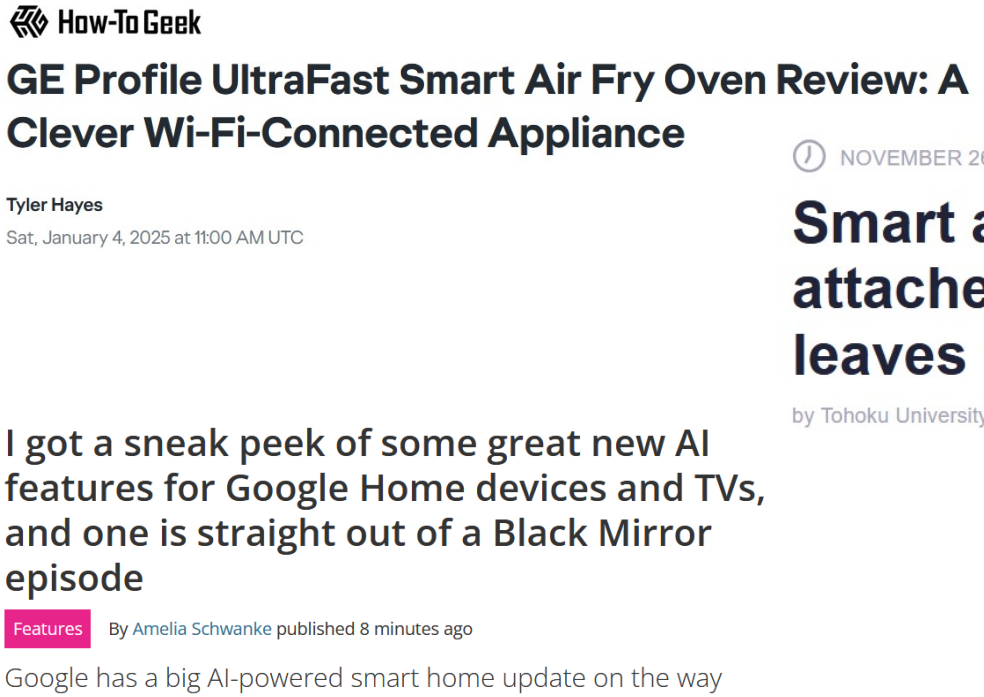

How-To Geek

GE Profile UltraFast Smart Air Fry Oven Review: A Clever Wi-Fi-Connected Appliance

Tyler Hayes

Sat, January 4, 2025 at 11:00 AM UTC

I got a sneak peek of some great new AI features for Google Home devices and TVs, and one is straight out of a Black Mirror episode

Features By Amelia Schwanke published 8 minutes ago

Google has a big AI-powered smart home update on the way

NOVEMBER 26, 2024

Editors' notes

Smart agriculture technology attaches directly to underside of leaves for monitoring plants

by Tohoku University

Cybersecurity studies how to detect, respond, and protect against attacks.

William E. Raymond, professional safecracker
Photographers: Harris & Ewing, 1937

https://www.loc.gov/resource/hec.33557/

Security is the protection of the confidentiality, integrity, and availability of systems and data.

"CIA Triad"

We'll cover (some of) the foundational approaches to AI

1. ∀x (Cat(x) → Mammal(x)) (All cats are mammals)
2. ∀x (Mammal(x) → Animal(x)) (All mammals are animals)
3. Cat(Tom) (Tom is a cat)

The Internet of Things connects objects, devices, sensors, and servers.

Google Nest Thermostat

GE Washer/Dryer

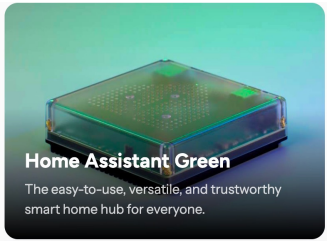

Hubitat

Google Nest Hub

Risk analysis is the heart of each topic of this class (and the field, if done right). Risk analysis is hard to get right.

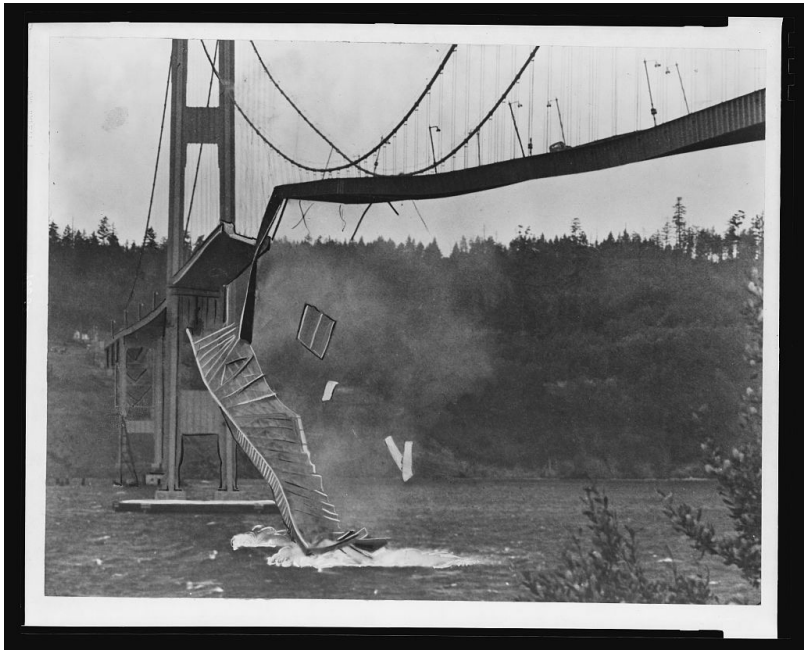

Tacoma Narrows Bridge Collapse (1940)

STS-51-L (Challenger) Disaster, January 28, 1986

Lehman Brothers bankruptcy, September 15, 2008

Q: What is an example of a computer security risk analysis that you did recently?

Risk management involves identifying threats, vulnerabilities, and impacts.

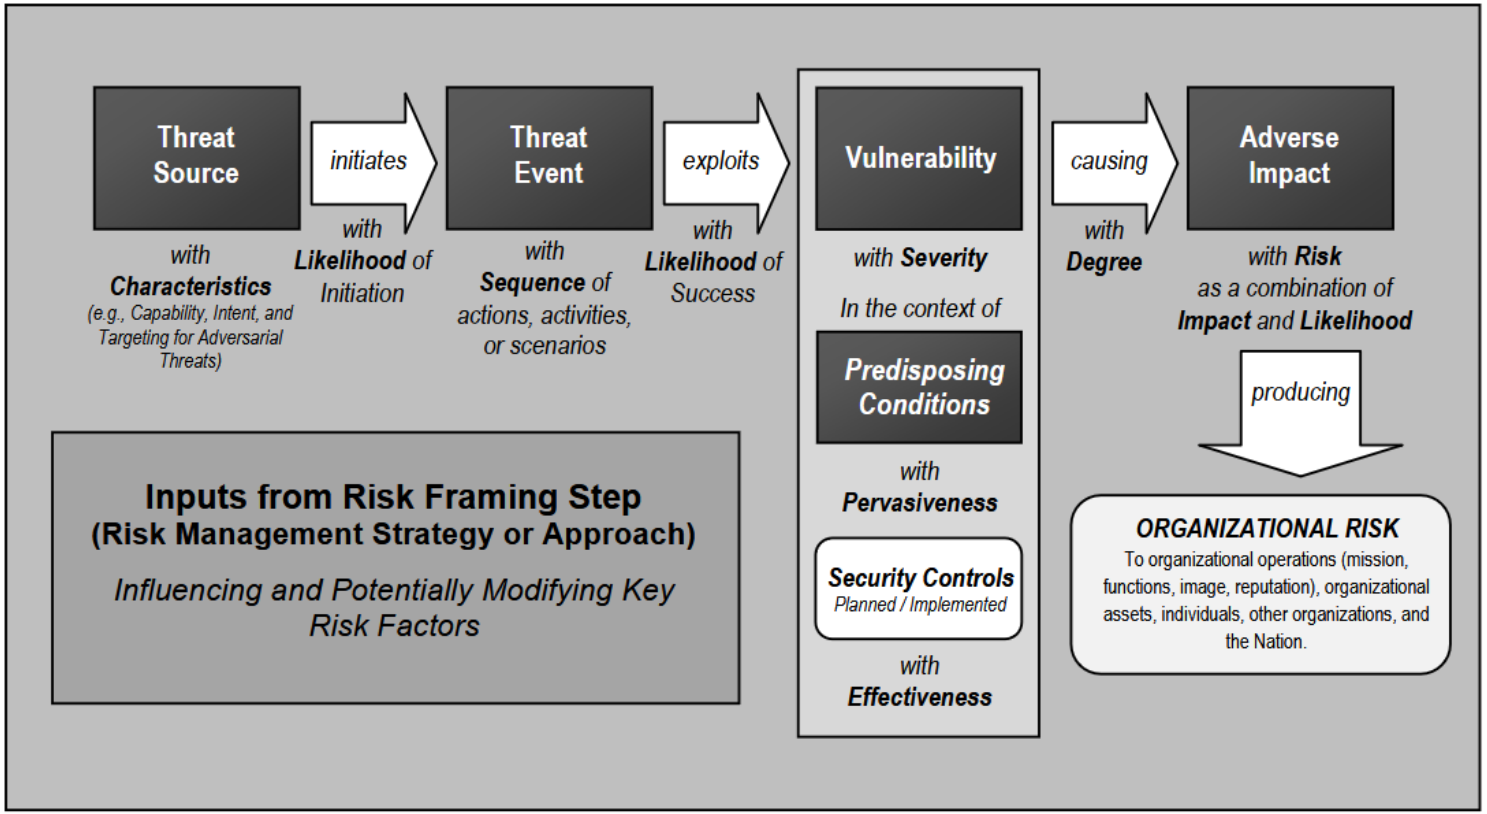

https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf

The RISKS Digest Archive

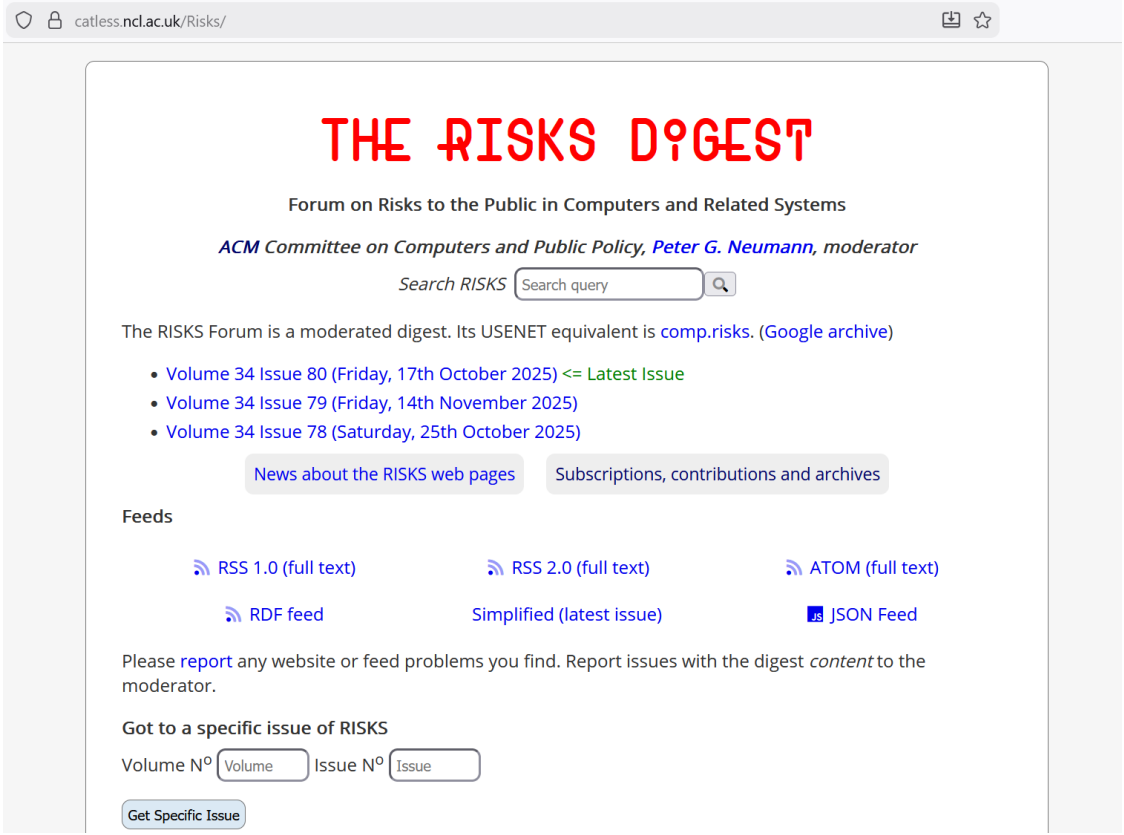

Generative AI Policy

AI Technologies Policy

Certain assignments in this course will permit or even encourage the use of generative artificial intelligence (GAI) tools such as ChatGPT. **The default is that such use is disallowed unless otherwise stated. Any such use must be appropriately acknowledged and cited**. It is each student's responsibility to assess the validity and applicability of any GAI output that is submitted; you bear the final responsibility. Violations of this policy will be considered academic misconduct. We draw your attention to the fact that different classes at Harvard could implement different AI policies, and it is the student's responsibility to conform to expectations for each course.

More Practical Generative AI Policy

It's okay to:

- Consult chatbots to ask questions about class topics. But be cautious about believing their answers.
- Ask chatbots for help debugging assignments (e.g., giving them the error messages).

It's not okay to:

- Use chatbots to write your comments to papers.

Course Management 2

The course has three main activities: lectures and participation; readings; and labs.

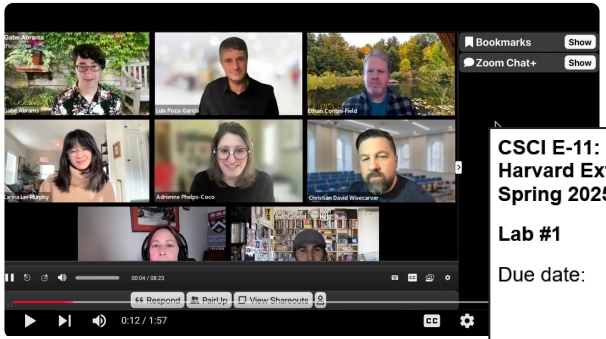

Lectures and participation

CSCI E-11: Artificial Intelligence, the Internet-of-Things, and Cybersecurity
Harvard Extension School
Spring 2025

Lab #1

Due date:

Introduction and Lab Objectives
Getting Started — First Read The Docs!
Create Your Account
Set a budget
Set up Multi-factor
Create a Micro Virtual Machine running AWS Linux
Tell Us About Your AWS Account
Install the git client and (optionally) an editor
Modify your system to retain logs for 6 months
Start Monitoring Your SSH logs
Get an A: Find our hack
What to turn in

Labs

"...no one can hack my mind": Comparing Expert and Non-Expert Security Practices

Iulia Ion, Google, iuliaion@google.com
Rob Reeder, Google, rreeder@google.com
Sunny Consolvo, Google, sconsolvo@google.com

ABSTRACT

The state of advice given to people today on how to stay safe online has plenty of room for improvement. Too many things are asked of them, which may be unrealistic, time consuming, or not really worth the effort. To improve the security advice, our community must find out what practices people use and what recommendations, if messaged well, are likely to bring the highest benefit while being realistic to ask of people. In this paper, we present the results of a study which aims to identify which practices people do that they consider most important at protecting their security online. We compare self-reported security practices of non-experts to those of security experts (i.e., participants who reported having five or more years of experience working in computer security). We report on the results of two online surveys—one with 231 security experts and one with 294 MTurk participants—on what the practices and attitudes of each group are. Our findings show a discrepancy between the security practices that experts and non-experts report taking. For instance, while experts most frequently report installing software updates, using two-factor authentication and using a password manager to stay safe online, non-experts report using antivirus software, visiting only known websites, and changing passwords frequently.

carefully considering the most worth-while advice to recommend is imperative. Even if users accept some responsibility for protecting their data [23, 43] and want to put in some effort [41], we should be thoughtful about what we ask them to do [20] and only offer advice that is effective and realistic to be followed.

Existing literature on giving good advice suggests that for recipients to follow it, the advice should be (a) useful, comprehensible and relevant, (b) effective at addressing the problem, (c) likely to be accomplished by the recipient, and (d) not possess too many limitations and drawbacks [34]. Therefore, to improve the state of security advice, we must assess which actions are most likely to be effective at protecting users, understand what users are likely and willing to do, and identify the potential challenges or inconveniences caused by following the advice. Furthermore, lessons from health advice in outreach interventions suggest that people will not initiate certain actions if they do not believe them to be effective [53]. Therefore, to learn how to best deliver the advice to users, we must also understand how users perceive its effectiveness and limitations.

In preliminary work, we surveyed security experts to identify what advice they would give non-tech-savvy users. The most frequently given pieces of advice were, in order of frequency: (1) keep systems and software up-to-date, (2) use unique passwords, (3) use strong passwords, (4) use two-factor authentication, (5) use antivirus software, and (6) use a password manager. In this paper, we report on results of a study which tries to identify what security advice users currently follow and how their attitudes and practices differ from those of security experts. To this end, we conducted a survey with 294 participants recruited from Amazon's Mechanical Turk crowdsourcing platform and another with 231 security experts recruited through an online blog. Our results help inform what important security advice users aren't following.

Our results show that expert participants considered keeping the operating system and applications up-to-date, using strong and unique passwords, turning on two-factor authentication, and using a password manager the most important things they do to stay safe online. Non-expert participants, however, considered using antivirus software, using strong passwords, changing passwords frequently, and visiting only trusted websites to be very effective, but admitted to delaying installation of software updates and expressed some lack of trust in password managers. We found that generally experts' security practices matched the advice they would give non-tech-savvy users, with a few exceptions. Experts recommended not clicking on links or opening emails from unknown people, yet they reported to do so at a higher rate than non-experts reported. Other security practices that non-experts considered very important, such as visiting only known websites, were not being followed by experts nor were they considered good security advice by experts.

1. INTRODUCTION

Frightening stories about cybersecurity incidents abound. The theft of millions of credit card numbers from a retail chain [10], a billion passwords from various websites [25], and a large set of nude celebrity photos [24] are just a few examples of stories that have been in the news lately.

In response to such security incidents, thousands of online articles and blog entries advise users what to do to stay safe online. Advice ranges from choosing a strong password [27] and having good security questions [38] to making email addresses unguessable [7] and entirely disabling photo backups in the cloud [27]. Besides such incident-related articles, many service providers, enterprises, and universities offer tips and training on how to stay safe online [2, 3, 17, 35].

If one hour of time from all US users is worth $2.5 billion [19],

Copyright is held by the author/owner. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee.
Symposium on Usable Privacy and Security (SOUPS) 2015, July 22–24, 2015, Ottawa, Canada.

USENIX Association 2015 Symposium on Usable Privacy and Security 327

Readings and quizzes

You must attend all class lectures (either live or recorded)

Online Synchronous (Live over Zoom): 6:30pm - 8:30 pm EST (GMT -5)
Online Asynchronous (Recorded with Immersive): Any time

You can pick which to do each week!

We also encourage you to self-organize “semi-synchronous” viewing sessions.

Ask Questions, Discuss the Class, or Get Help on Ed Discussion Boards

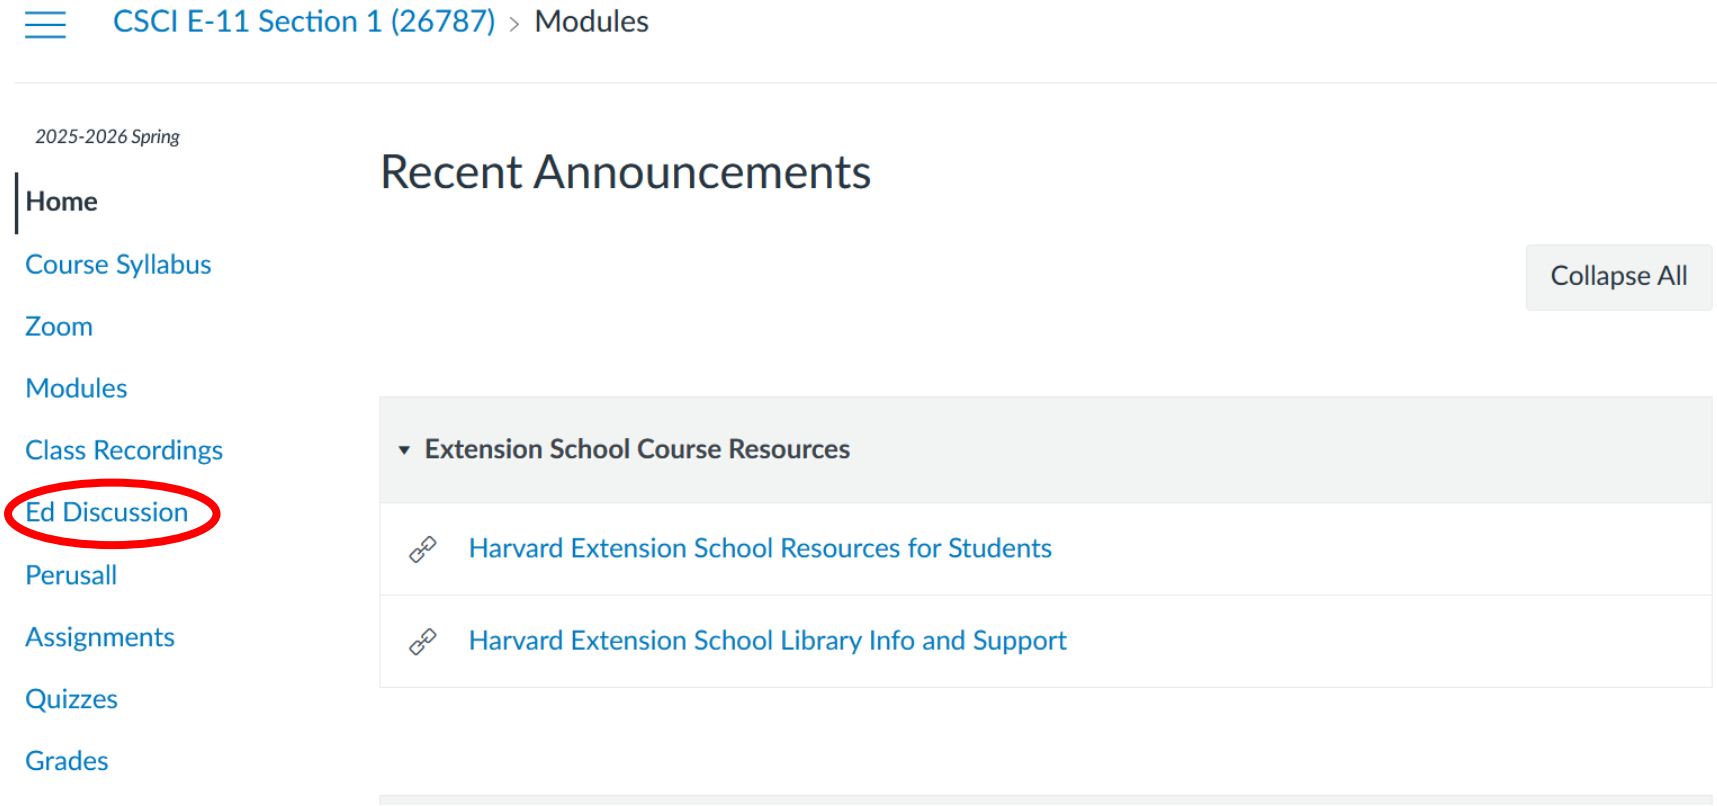

Ed Discussion board is the best place to get help and to participate in discussions.

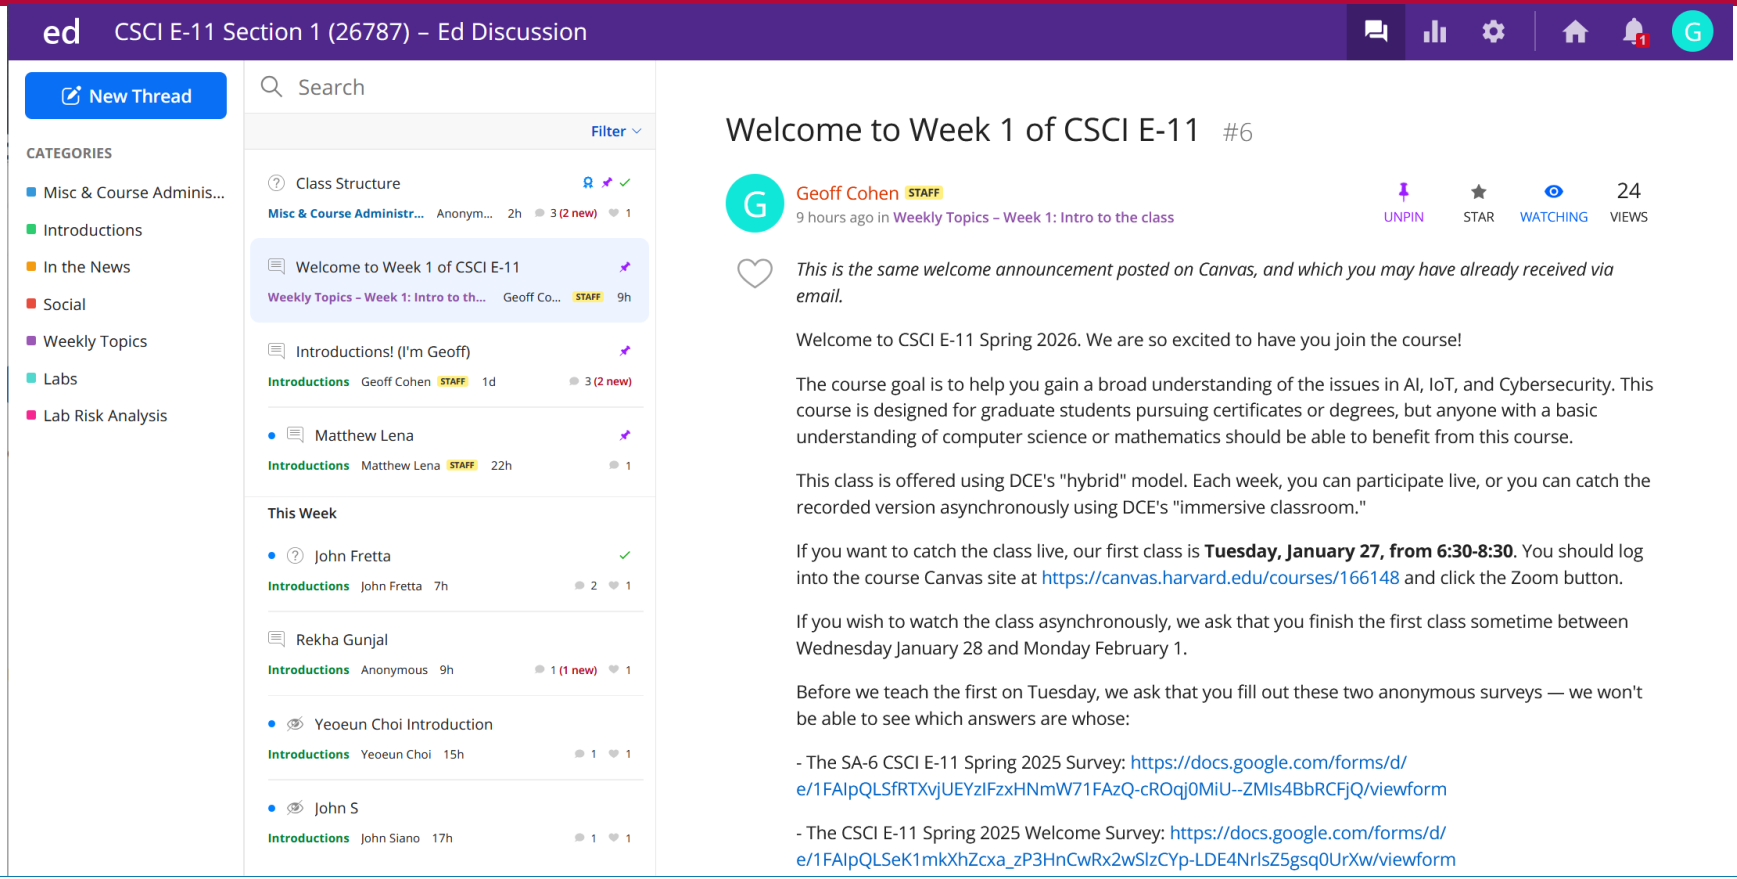

Grading

Weekly Quizzes on readings:

- 14 readings @ 2 points each = 28 points

Labs:

- 8 labs @ 5 points each = 40 points
  - *The lab assignment provides all of the information you need for a passing grade.*
  - *To get an "A" you need to figure something out.*
  - *For some labs, we give you ideas of how to go further — please feel free to do more and share what you have done in the discussion forums!*
- 8 lab quizzes @ 2 points each = 16 points

Class Participation: 15%

- *Attend each lecture.*
- *Participate in Perusall, the immersive classroom, and on the discussion boards.*
- *Aim for 2-4 interactions each week.*

Participation

- Come to class! (Live or recorded)
  - Comment in the Zoom chat (works for the recorded session, too)
- Comment on the weekly reading on Perusall
  - (At least) one substantive comment
  - (At least) one substantive reply to someone else's comment
- Participate on the Ed Discussion board

- We'll provide some more guidance next week.

Late Work Policy

Submit your work on time to receive credit

*If you have a school-approved accommodation in place, make sure that the instructors know, and ask for schedule changes as soon as possible.

Stay up to date — We will send out announcements using Canvas

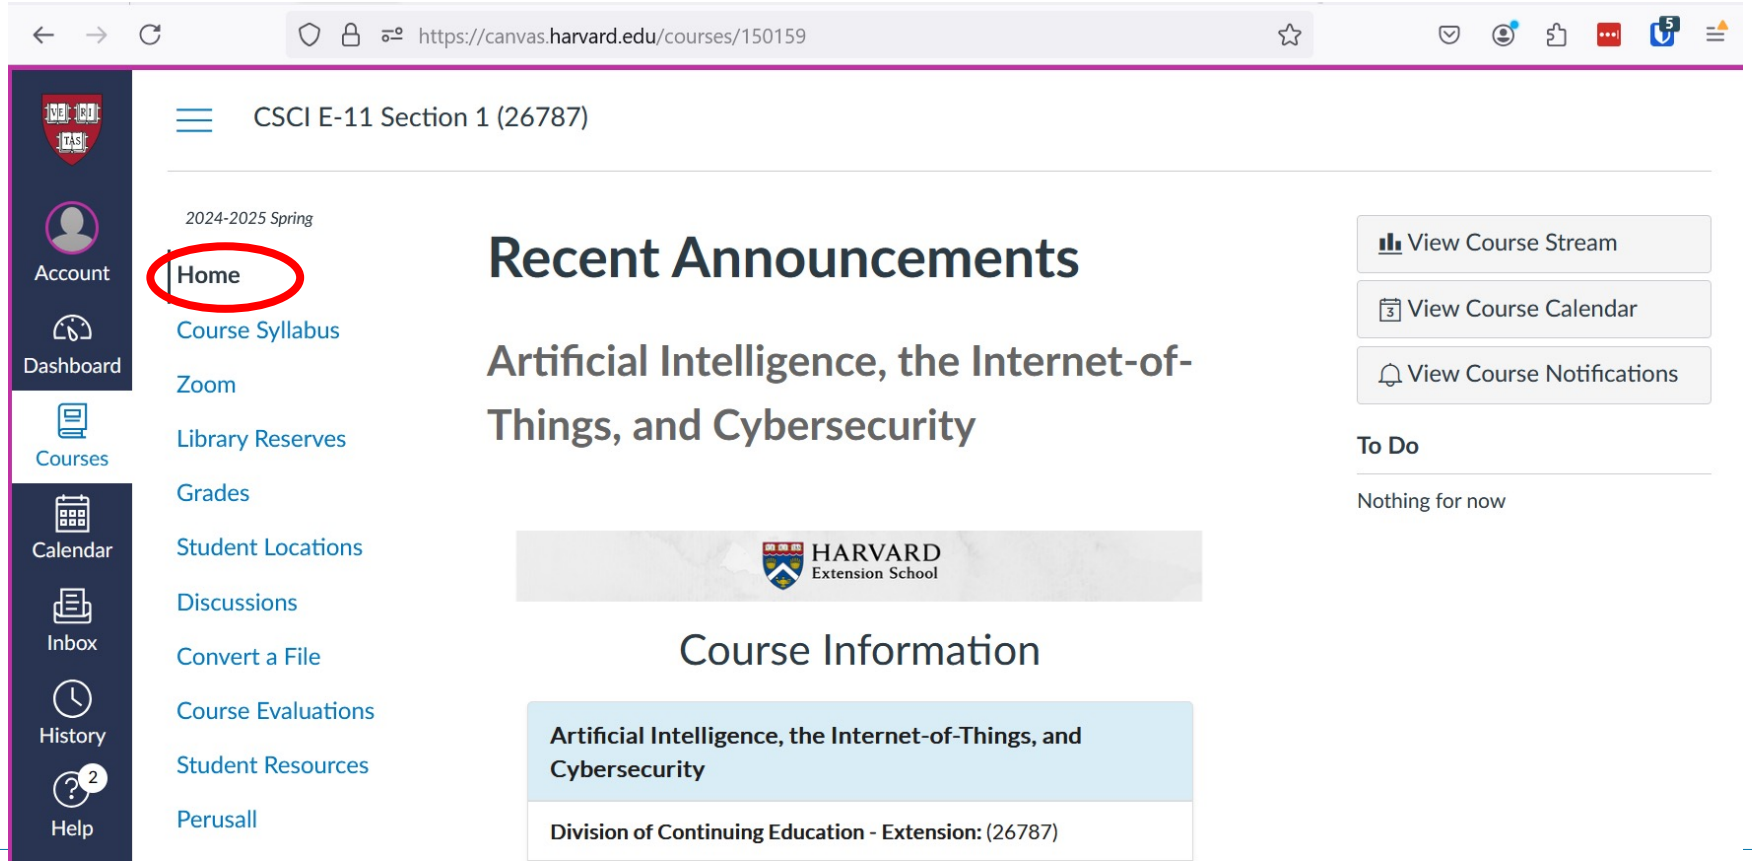

Use Immersive Video to watch recorded lectures.

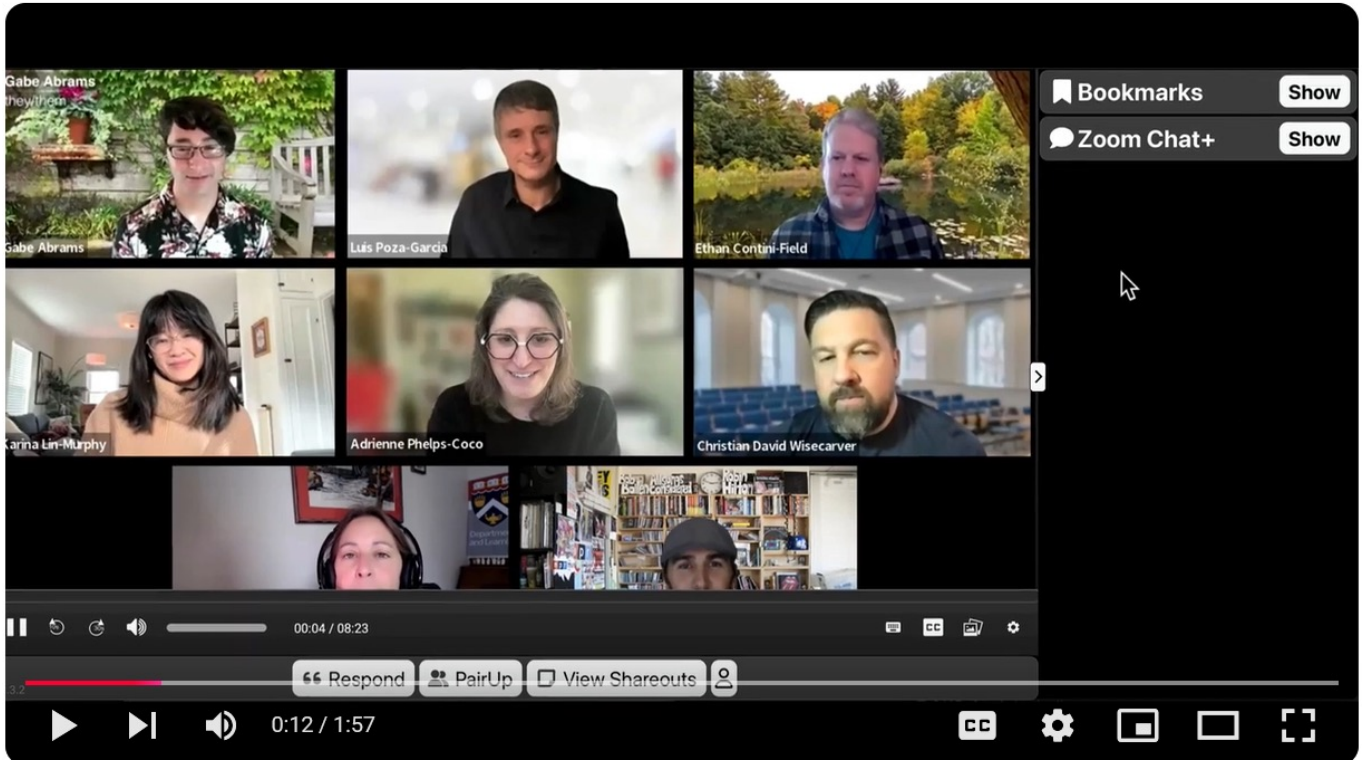

Use the “Study Lounge” to Zoom chat with other students who are online at the same time that you are!

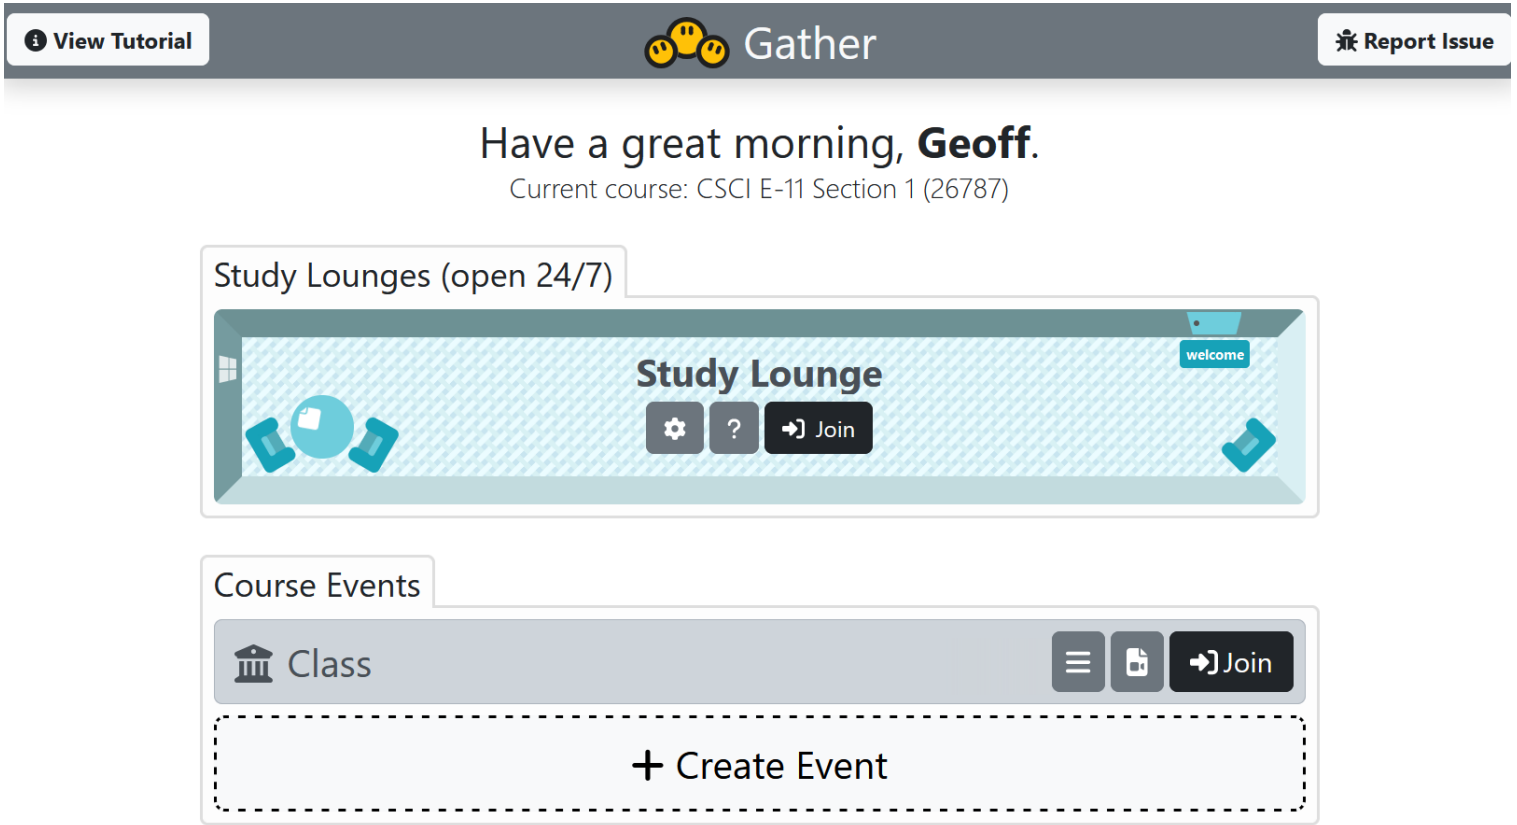

Harvard's Zoom Code of Conduct

Code of Conduct:

This space has been created to allow for interactions between students outside course meeting times. In joining this room, you are taking responsibility for following DCE's code of conduct and for creating a positive and supportive environment. In order to keep this room safe and welcoming for everyone, you must agree to the following before joining:

- You will abide by DCE standards of academic integrity and the collaboration policies set forth by your class
- You will use appropriate language and treat all classmates with respect
- You will dress in clothing you would wear to class if you attended in person
- You will refrain from driving a vehicle while in the Study Lounge

The class readings are on Perusall

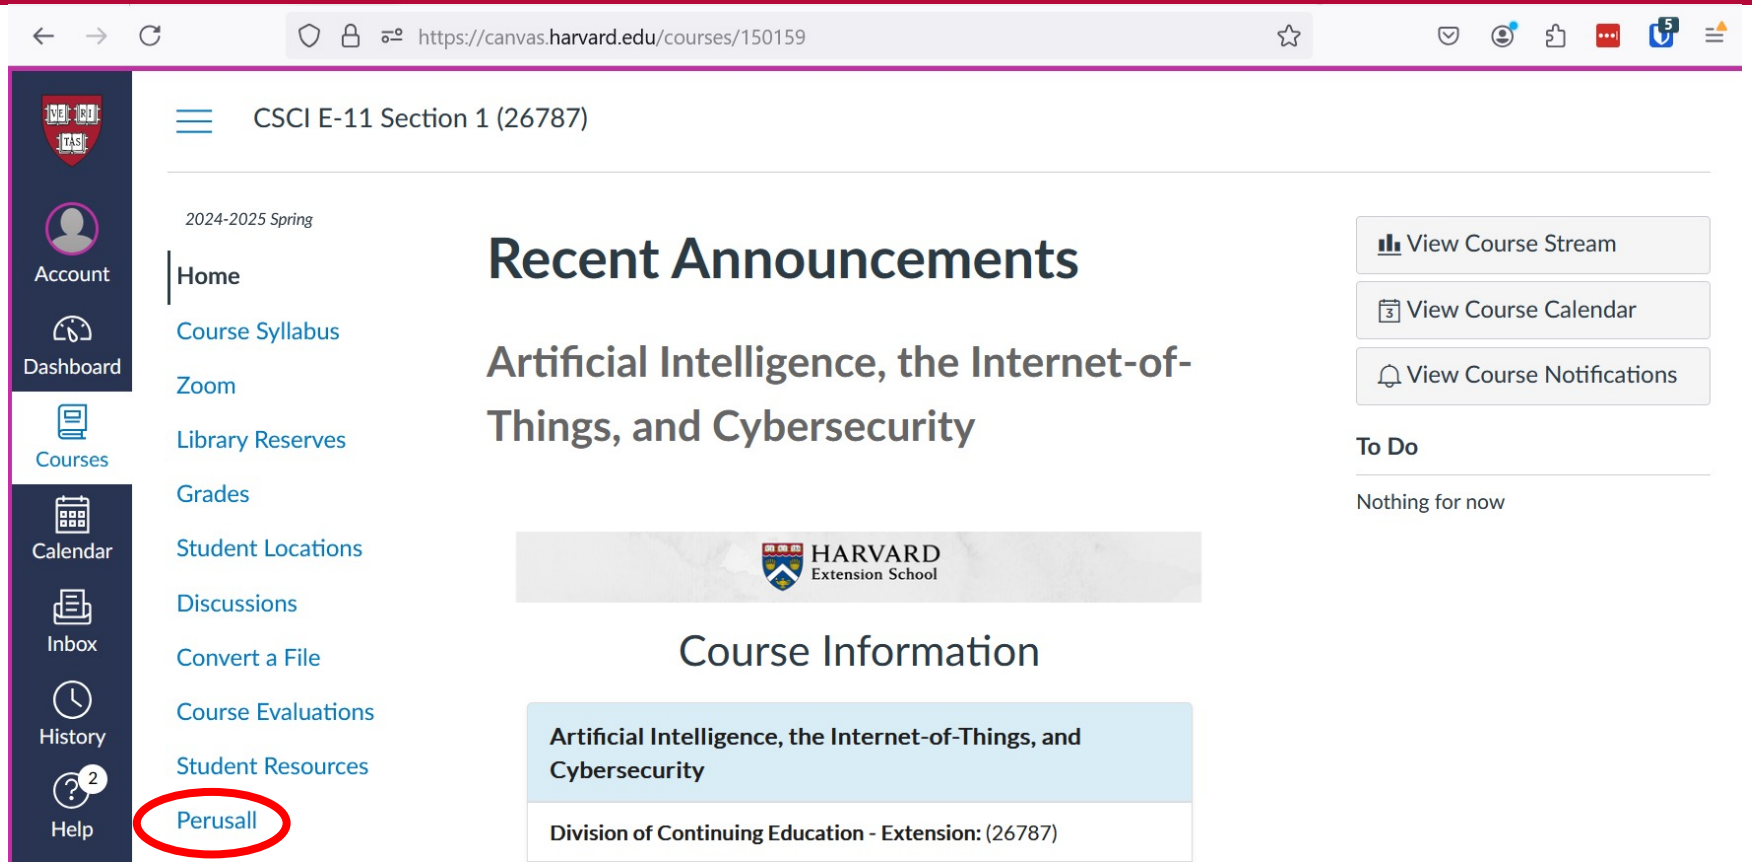

Use Perusall to comment on or ask questions about the papers.

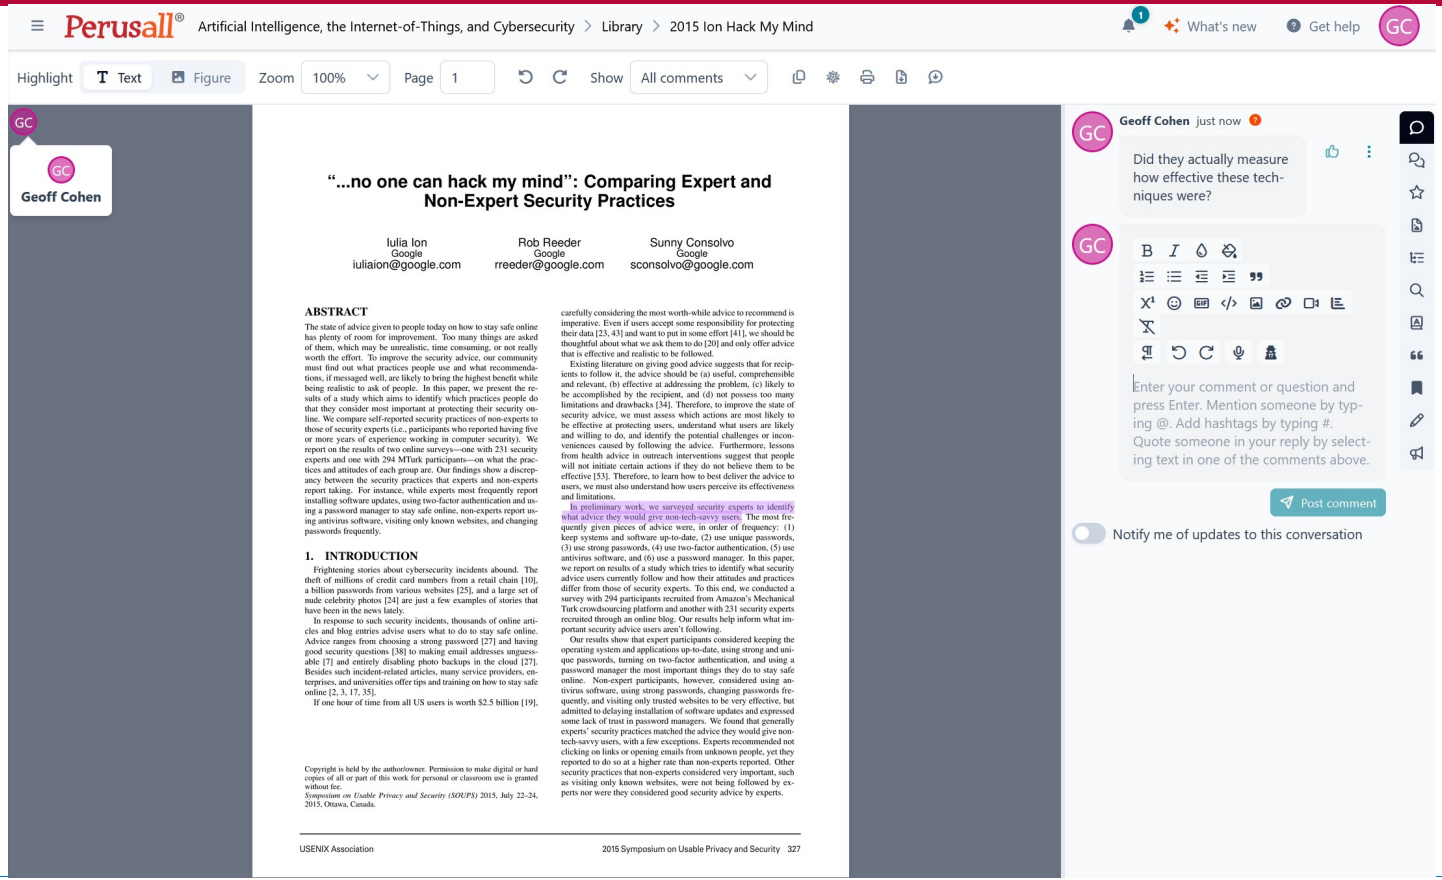

After you read the paper, answer the Quiz on Canvas

"...no one can hack my mind": Comparing Expert and Non-Expert Security Practices

Iulia Ion
Google
iuliaion@google.com

Rob Reeder
Google
rreeder@google.com

Sunny Consolvo
Google
sconsolvo@google.com

ABSTRACT

The state of advice given to people today on how to stay safe online has plenty of room for improvement. Too many things are asked of them, which may be unrealistic, time consuming, or not really worth the effort. To improve the security advice, our community must find out what practices people use and what recommendations, if messaged well, are likely to bring the highest benefit while being realistic to ask of people. In this paper, we present the results of a study which aims to identify which practices people do that they consider most important at protecting their security online. We compare self-reported security practices of non-experts to those of security experts (i.e., participants who reported having five or more years of experience working in computer security). We report on the results of two online surveys—one with 231 security experts and one with 294 MTurk participants—on what the practices and attitudes of each group are. Our findings show a discrepancy between the security practices that experts and non-experts report taking. For instance, while experts most frequently report installing software updates, using two-factor authentication and using a password manager to stay safe online, non-experts report using antivirus software, visiting only known websites, and changing passwords frequently.

1. INTRODUCTION

Frightening stories about cybersecurity incidents abound. The theft of millions of credit card numbers from a retail chain [10], a billion passwords from various websites [25], and a large set of nude celebrity photos [24] are just a few examples of stories that have been in the news lately.

In response to such security incidents, thousands of online articles and blog entries advise users what to do to stay safe online. Advice ranges from choosing a strong password [27] and having good security questions [38] to making email addresses unguessable [7] and entirely disabling photo backups in the cloud [27]. Besides such incident-related articles, many service providers, enterprises, and universities offer tips and training on how to stay safe online [2, 3, 17, 35].

If one hour of time from all US users is worth $2.5 billion [19], carefully considering the most worth-while advice to recommend is imperative. Even if users accept some responsibility for protecting their data [23, 43] and want to put in some effort [41], we should be thoughtful about what we ask them to do [20] and only offer advice that is effective and realistic to be followed.

Existing literature on giving good advice suggests that for recipients to follow it, the advice should be (a) useful, comprehensible and relevant, (b) effective at addressing the problem, (c) likely to be accomplished by the recipient, and (d) not possess too many limitations and drawbacks [34]. Therefore, to improve the state of security advice, we must assess which actions are most likely to be effective at protecting users, understand what users are likely and willing to do, and identify the potential challenges or inconveniences caused by following the advice. Furthermore, lessons from health advice in outreach interventions suggest that people will not initiate certain actions if they do not believe them to be effective [53]. Therefore, to learn how to best deliver the advice to users, we must also understand how users perceive its effectiveness and limitations.

In preliminary work, we surveyed security experts to identify what advice they would give non-tech-savvy users. The most frequently given pieces of advice were, in order of frequency: (1) keep systems and software up-to-date, (2) use unique passwords, (3) use strong passwords, (4) use two-factor authentication, (5) use antivirus software, and (6) use a password manager. In this paper, we report on results of a study which tries to identify what security advice users currently follow and how their attitudes and practices differ from those of security experts. To this end, we conducted a survey with 294 participants recruited from Amazon's Mechanical Turk crowdsourcing platform and another with 231 security experts recruited through an online blog. Our results help inform what important security advice users aren't following.

Our results show that expert participants considered keeping the operating system and applications up-to-date, using strong and unique passwords, turning on two-factor authentication, and using a password manager the most important things they do to stay safe online. Non-expert participants, however, considered using antivirus software, using strong passwords, changing passwords frequently, and visiting only trusted websites to be very effective, but admitted to delaying installation of software updates and expressed some lack of trust in password managers. We found that generally experts' security practices matched the advice they would give non-tech-savvy users, with a few exceptions. Experts recommended not clicking on links or opening emails from unknown people, yet they reported to do so at a higher rate than non-experts reported. Other security practices that non-experts considered very important, such as visiting only known websites, were not being followed by experts nor were they considered good security advice by experts.

Symposium on Usable Privacy and Security (SOUPS) 2015, July 22–24, 2015, Ottawa, Canada.

Quiz Instructions

Please answer the following multiple-choice questions based on the paper, "No One Can Hack My Mind: Comparing Expert and Non-Expert Security Practices" by Iulia Ion, Rob Reeder, and Sunny Consolvo.

Question 1 — 1 pts

What are characteristics of good advice? Select all of the best answers.

- ☐ It should offer the best possible approach.
- ☐ It should be likely to be accomplished by the recipient.
- ☐ It should not contain any limitations or drawbacks.
- ☐ It should be effective at addressing the problem.

Question 2 — 1 pts

According to the paper, what's an example of a security practice that is far more likely to be followed by an expert compared to a non-expert?

- ☐ Don't ever share personal information online.
- ☐ Visit only known websites.
- ☐ Use anti-virus software.
- ☐ Use two-factor authentication.

Labs

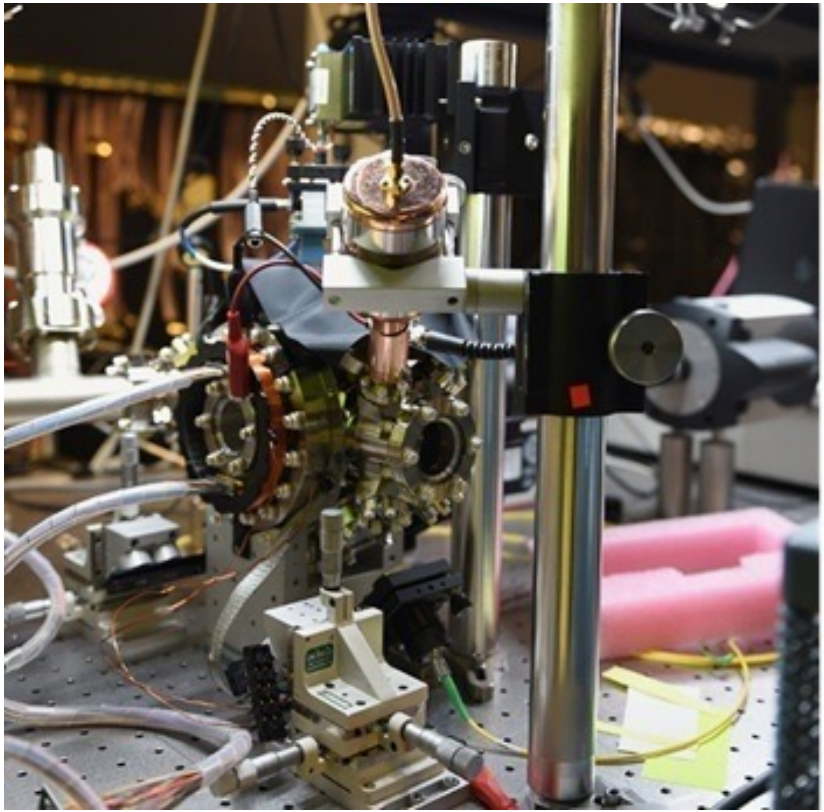

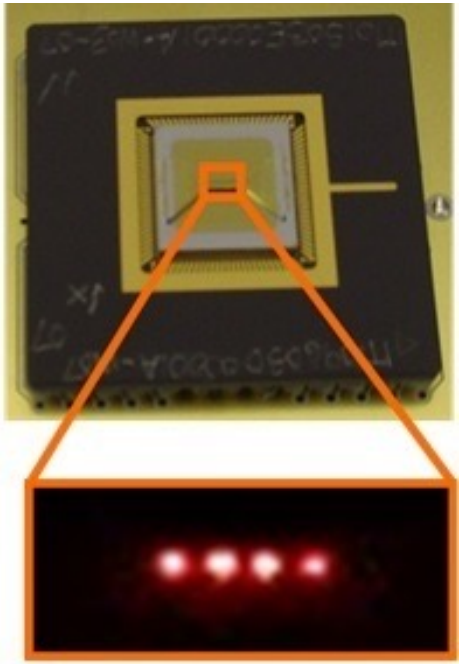

YB+ Ion Trapping (Wright-Patterson AFB, 2017)

Lab #0: Please purchase a MEMENTO from AdaFruit

Sensors / Cameras / MEMENTO - Python Programmable DIY Camera - Bare Board

You last purchased this item on December 23, 2025.
View this order

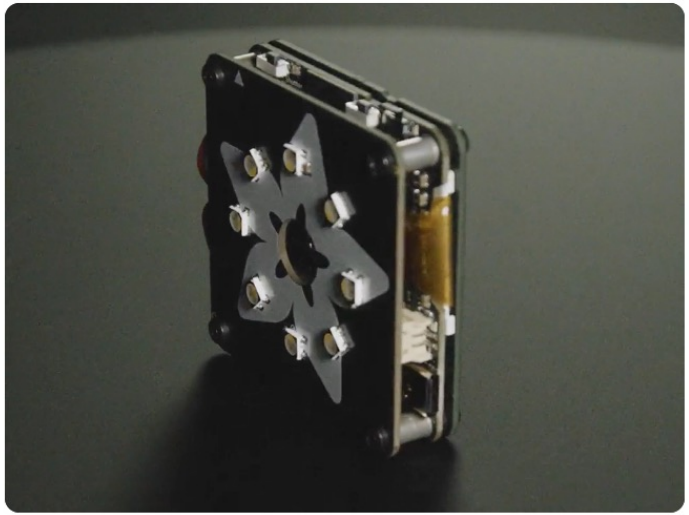

MEMENTO - Python Programmable DIY Camera - Bare Board

Product ID: 5420

$34.95
88 in stock

1 Add to Cart

Also include 1 x Square Maker-Friendly Zipper Case - Purple ($2.50)

Add to Wishlist

Description

Technical Details

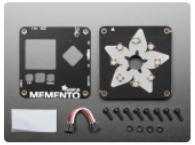

Adafruit MEMENTO Camera Enclosure & Hardware Kit

PRODUCT ID: 5843

Add to Cart

$9.95
32 in stock

Also get...

You last purchased this item on December 23, 2025.
View this order

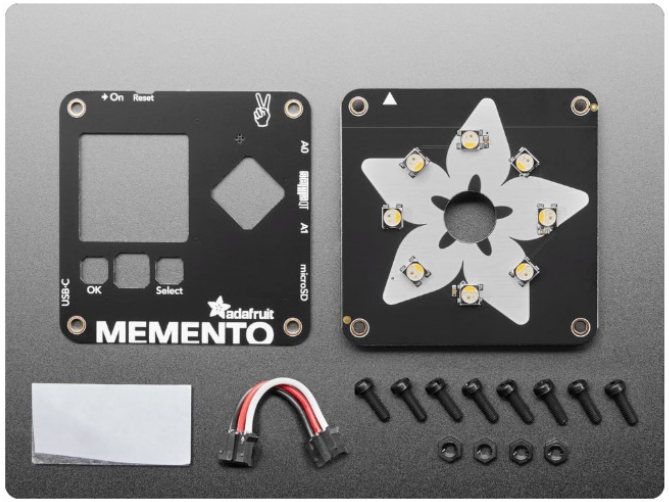

Adafruit MEMENTO Camera Enclosure & Hardware Kit

Product ID: 5843

$9.95

32 in stock

1 Add to Cart

Qty	Discount
1-9	$9.95
10-99	$8.96
100+	$7.96

Add to Wishlist

Description

Technical Details

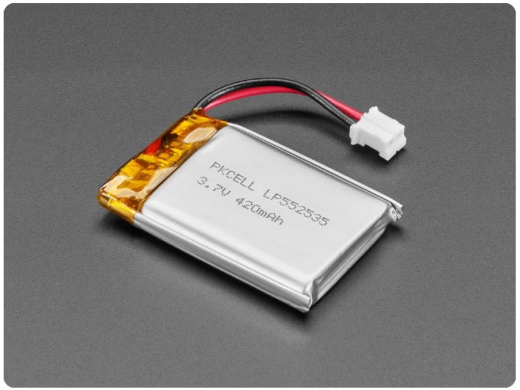

Lithium Ion Polymer Battery with Short Cable - 3.7V 420mAh

Product ID: 4236

$6.95

There are multiple versions of this item. Please select one from the options below:

3.7V 420mAh with Short Cable

60 in stock

1 Add to Cart

Qty	Discount
1-9	$6.95
10-99	$6.26
100+	$5.56

Add to Wishlist

You will create an AWS account for this course in Lab #1.
Please do this — even if you already have one!

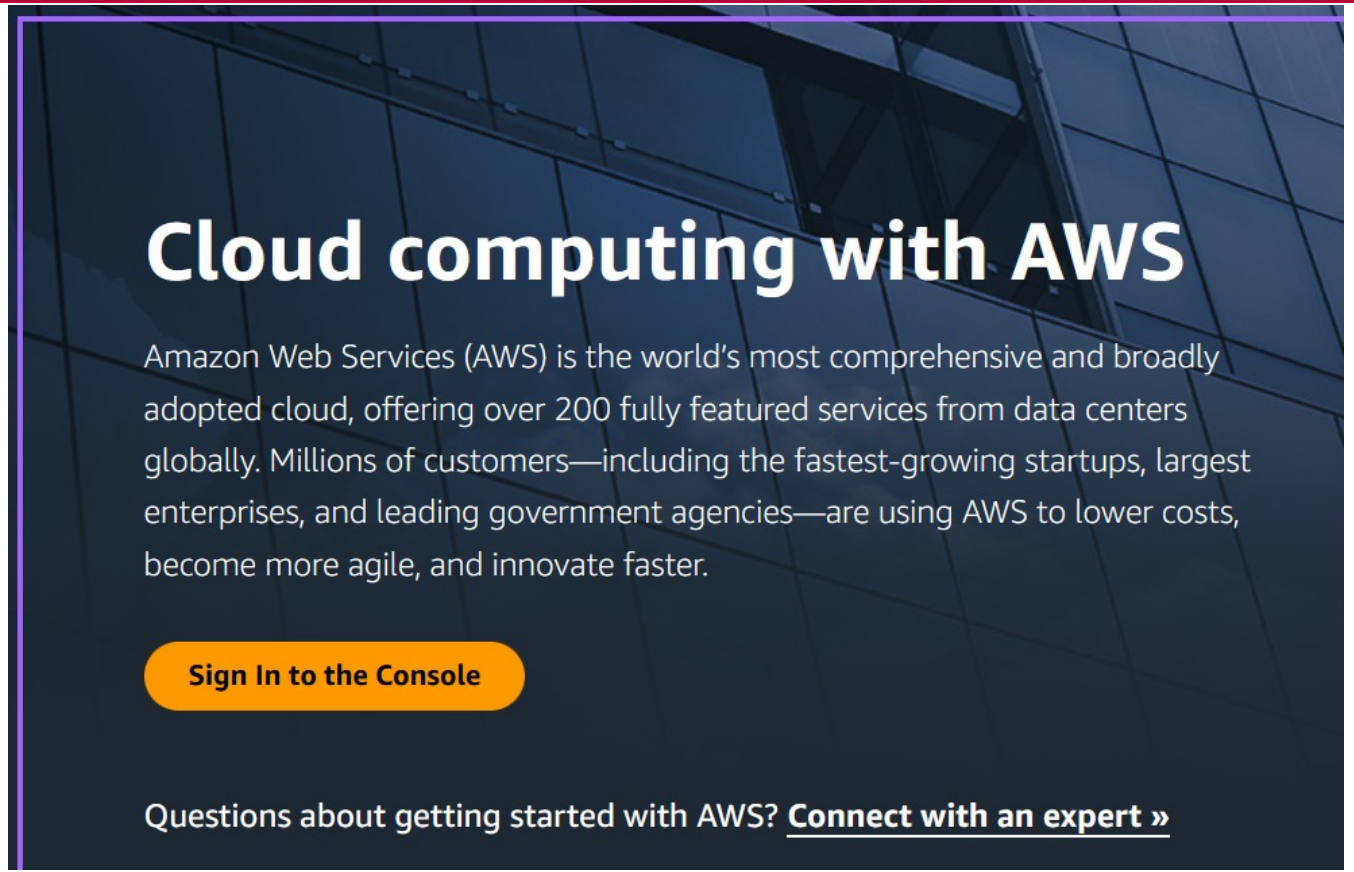

Lab Locations

All of the labs are linked from our GitHub README file:

- https://github.com/Harvard-CSCI-E-11/spring26

To use the labs, you must access the lab dashboard at:

- https://csci-e-11.org/

You need to Harvard Key to access the labs!

- It's the same Harvard Key you use to access Canvas!

Once you log in, write down your "course key"

The dashboard…

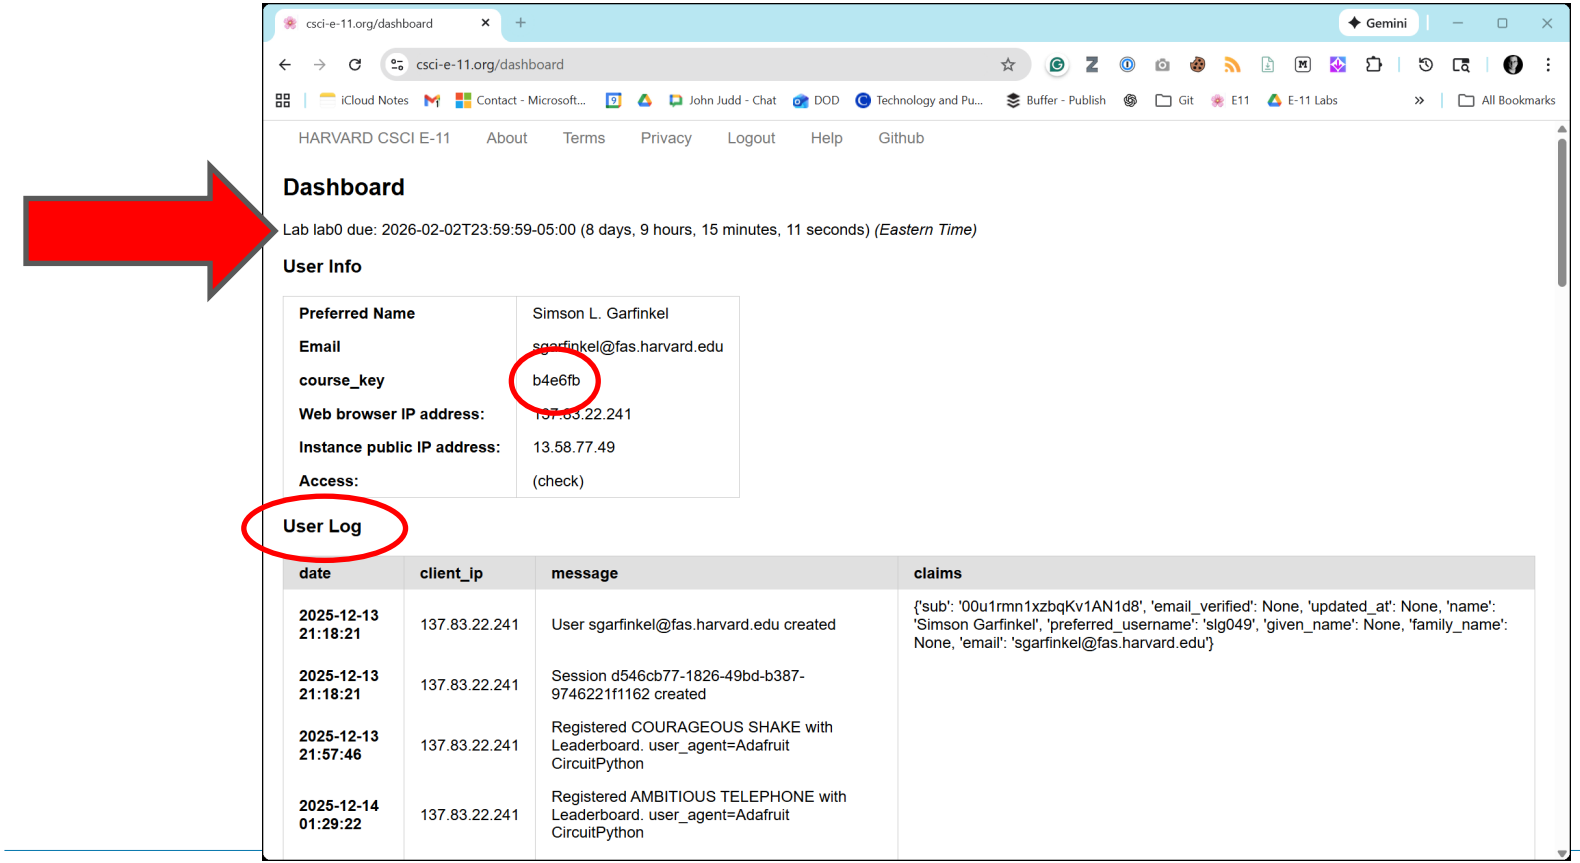

Lab flow…

1. Read the lab description on Google Docs
2. If something isn't clear, post in the lab discussion on the Ed Board
    1. Other students can help you!
    2. The TAs can help you!
3. Log into your AWS EC2 Instance with "ssh" or the AWS web-based "connect"
4. `cd $HOME/spring26/{labN}`
5. Work on the lab…
6. `$ e11 check {labN}` --- *Runs the grader locally*
7. `$ e11 grade {labN}` --- *Regrade as often as you want; highest grade wins!*

We will automatically grade all students that don't have a 5/5 at day -4, -3, -2, -1 and 0!

Please let us know if you have problems!

This is the first time we've used the dashboard and the automatic grader!

(Simson wrote it last fall…)

You will be using git – the distributed source code control system.

Lab #1
Creating a Server in Amazon Web Services

Create an AWS Account

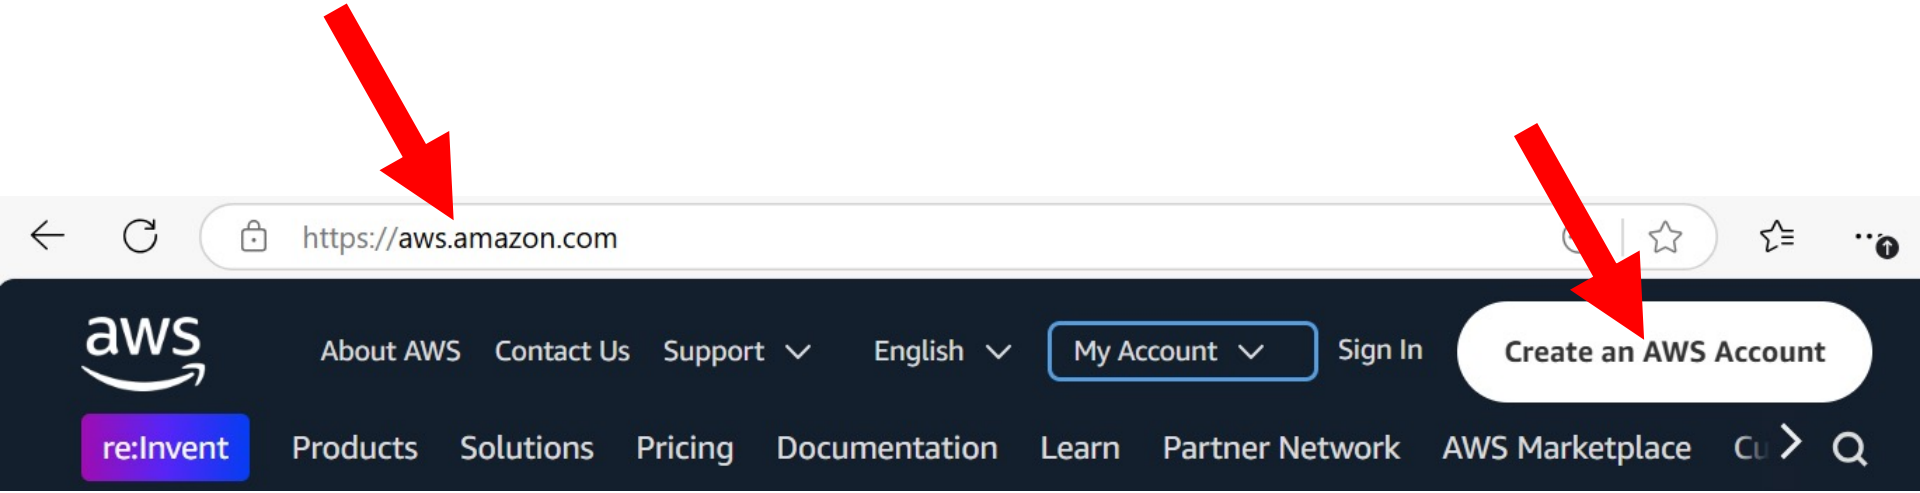

We will be using the AWS Free Tier in this course.

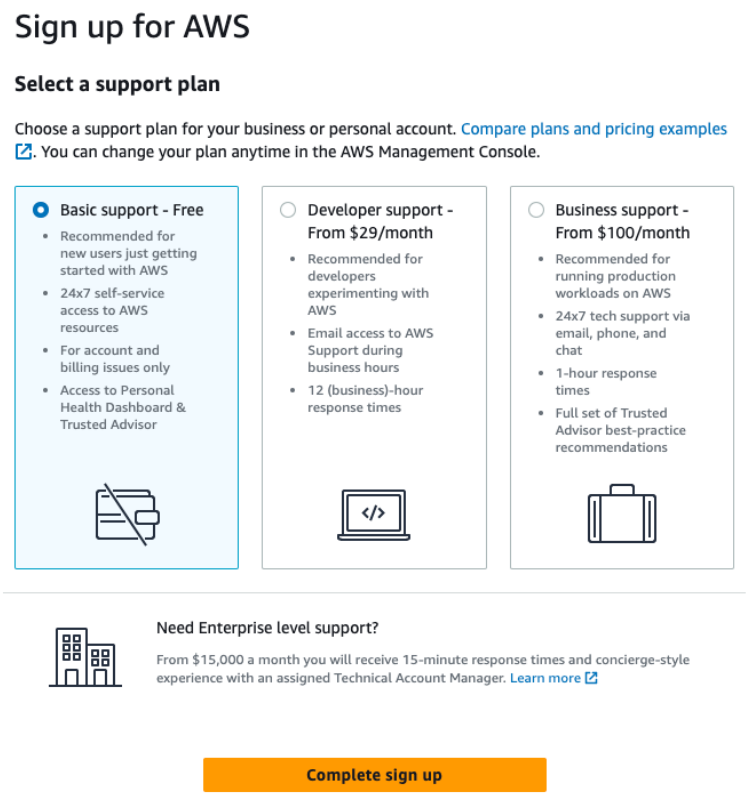

Be sure to set a budget so you are warned if you start spending money!

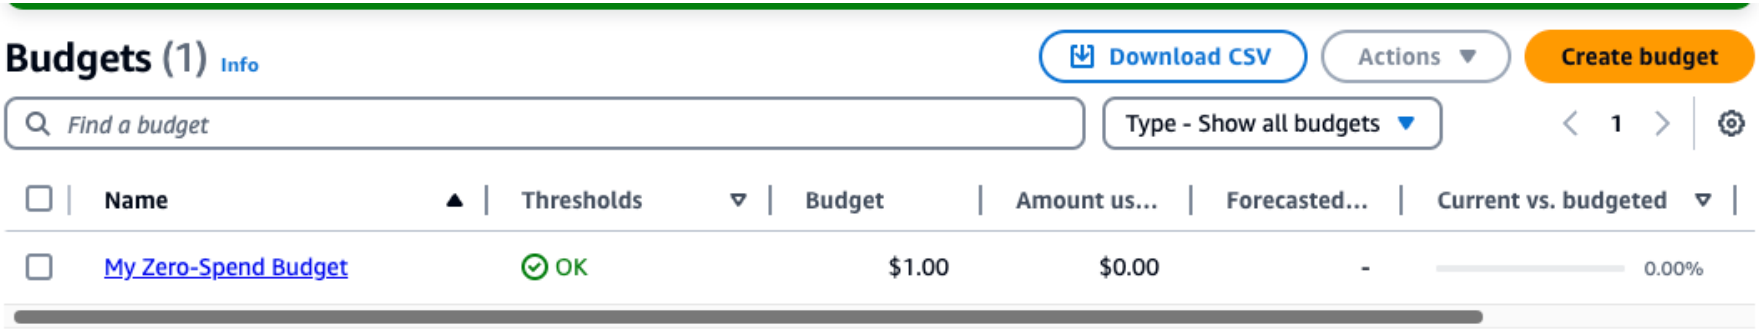

Create an “Instance” —a Virtual Server — in Amazon’s Elastic Compute Cloud.

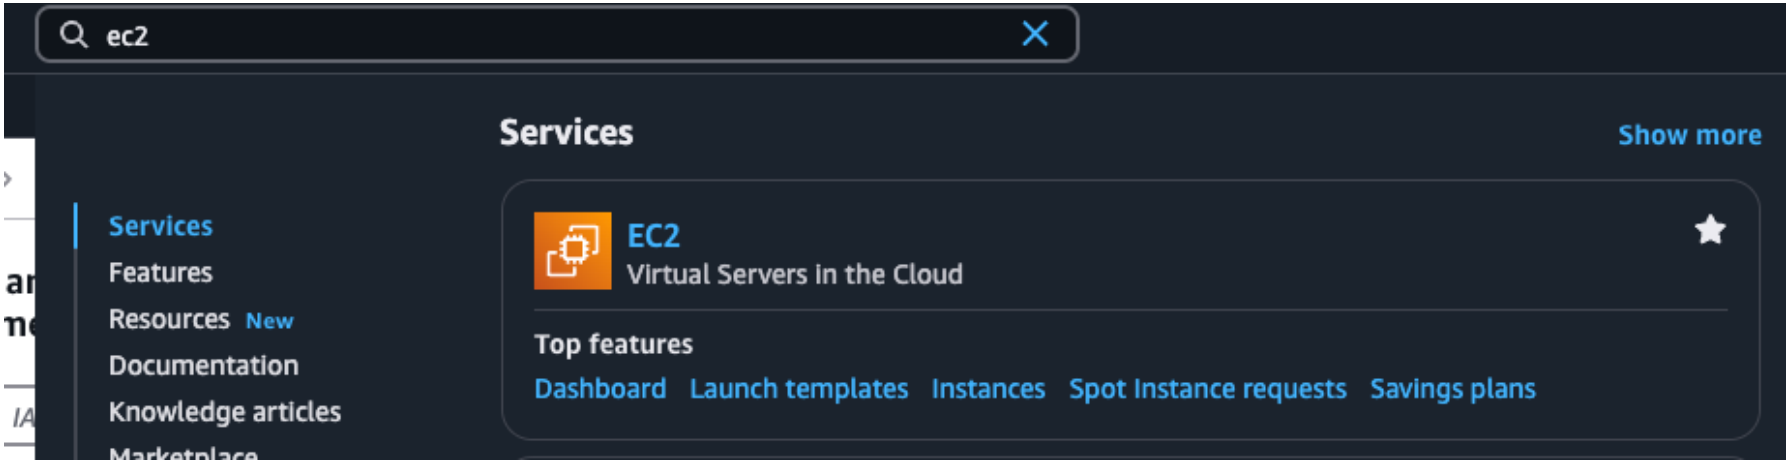

Log in to the Amazon Linux Command Line

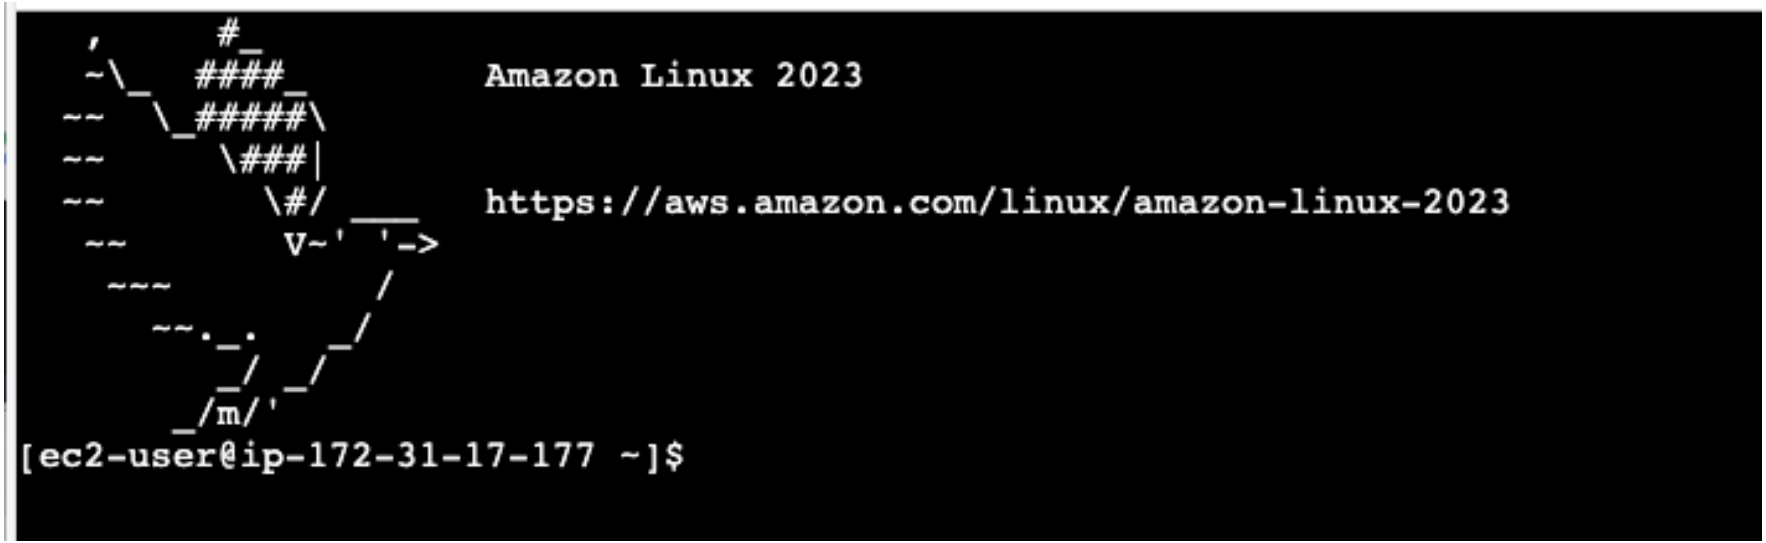

We have created a system called “e11” to help you in this course. It will give you information each time you log in...

```
Welcome to Ubuntu 24.04.2 LTS (GNU/Linux 6.14.0-1010-aws aarch64)

 * Documentation:  https://help.ubuntu.com
 * Management:     https://landscape.canonical.com
 * Support:        https://ubuntu.com/pro

 System information as of Fri Dec 26 20:13:14 UTC 2025

  System load:  0.0               Temperature:           -273.1 C
  Usage of /:   68.1% of 6.71GB   Processes:             149
  Memory usage: 47%               Users logged in:       0
  Swap usage:   0%                IPv4 address for ens5: 172.31.8.79

 * Ubuntu Pro delivers the most comprehensive open source security and
   compliance features.

   https://ubuntu.com/aws/pro

Expanded Security Maintenance for Applications is not enabled.

62 updates can be applied immediately.
To see these additional updates run: apt list --upgradable

Enable ESM Apps to receive additional future security updates.
See https://ubuntu.com/esm or run: sudo pro status


*** System restart required ***
Last login: Wed Dec 24 22:21:40 2025 from 137.83.22.241
E11 local version: 0.2.2
E11 server version: 0.2.2 (deployed 2025-12-26T15:03:16-0500)
uCSCI E-11 Course admin and grader HAVE ACCESS to this instance.
               total        used        free      shared  buff/cache   available
Mem:           909Mi       421Mi        90Mi       2.2Mi       495Mi       487Mi
● lab6.service - Gunicorn instance to serve Flask application for lab 6
     Loaded: loaded (/etc/systemd/system/lab6.service; disabled; preset: enabled)
     Active: active (running) since Wed 2025-12-24 20:48:53 UTC; 1 day 23h ago
   Main PID: 18066 (gunicorn)
      Tasks: 5 (limit: 1080)
     Memory: 136.3M (peak: 141.9M)
        CPU: 15min 35.906s
     CGroup: /system.slice/lab6.service
             ├─18066 /home/ubuntu/spring26/lab6/.venv/bin/python /home/ubuntu/spring26/lab6/.venv/bin/gunicorn -w 2 -b 127.0.0.1:8006 --reload server.main:a…
             ├─18598 /home/ubuntu/spring26/lab6/.venv/bin/python /home/ubuntu/spring26/lab6/.venv/bin/gunicorn -w 2 -b 127.0.0.1:8006 --reload server.main:a…
             └─18599 /home/ubuntu/spring26/lab6/.venv/bin/python /home/ubuntu/spring26/lab6/.venv/bin/gunicorn -w 2 -b 127.0.0.1:8006 --reload server.main:a…
Warning: The unit file, source configuration file or drop-ins of lab6.service changed on disk. Run 'systemctl daemon-reload' to reload units.
```

```
E11 local version: 0.2.2
E11 server version: 0.2.2 (deployed 2025-12-26T15:03:16-0500)
uCSCI E-11 Course admin and grader HAVE ACCESS to this instance.
               total        used        free      shared  buff/cache   available
Mem:           909Mi       421Mi        90Mi       2.2Mi       495Mi       487Mi
● lab6.service - Gunicorn instance to serve Flask application for lab 6
     Loaded: loaded (/etc/systemd/system/lab6.service; disabled; preset: enabled)
     Active: active (running) since Wed 2025-12-24 20:48:53 UTC; 1 day 23h ago
   Main PID: 18066 (gunicorn)
      Tasks: 5 (limit: 1080)
     Memory: 136.3M (peak: 141.9M)
        CPU: 15min 35.906s
     CGroup: /system.slice/lab6.service
             ├─18066 /home/ubuntu/spring26/lab6/.venv/bin/python /home/ubuntu/spring26/lab6/.venv/bin/gunicorn -w 2 -b 127.0.0.1:8006 --reload server.main:a…
             ├─18598 /home/ubuntu/spring26/lab6/.venv/bin/python /home/ubuntu/spring26/lab6/.venv/bin/gunicorn -w 2 -b 127.0.0.1:8006 --reload server.main:a…
             └─18599 /home/ubuntu/spring26/lab6/.venv/bin/python /home/ubuntu/spring26/lab6/.venv/bin/gunicorn -w 2 -b 127.0.0.1:8006 --reload server.main:a…
Warning: The unit file, source configuration file or drop-ins of lab6.service changed on disk. Run 'systemctl daemon-reload' to reload units.
```

If we change e11 during the course, you can type `e11 update` to get the new version...

```
ubuntu@simsong20acm:~$ e11 update
$ cd /home/ubuntu/spring26
$ git stash
Saved working directory and index state WIP on main: 2189320 Merge pull request #138 from Harvard-CSCI-E-11/test-get_parser
$ git pull
remote: Enumerating objects: 99, done.
remote: Counting objects: 100% (49/49), done.
remote: Compressing objects: 100% (16/16), done.
remote: Total 99 (delta 34), reused 40 (delta 33), pack-reused 50 (from 1)
Unpacking objects: 100% (99/99), 70.19 KiB | 1.80 MiB/s, done.
From https://github.com/Harvard-CSCI-E-11/spring26
   2189320..a5f9300  main            -> origin/main
 * [new branch]      simsong-patch-1 -> origin/simsong-patch-1
Updating 2189320..a5f9300
Fast-forward
 etc/e11-cli/.gitignore                                          |   3 +
 etc/e11-cli/Makefile                                            |  15 +++
 etc/e11-cli/e11/e11_common.py                                   |  22 +++-
 etc/e11-cli/e11/e11core/context.py                              |   8 +-
 etc/e11-cli/e11/lab_tests/lab7_test.py                          |  41 +++++++
 etc/e11-cli/e11/lab_tests/lab8_test.py                          | 195 ++++++++++++++++++++++++++++++++++
 etc/e11-cli/lambda-home/src/home_app/templates/dashboard.html   |   8 +-
```

Extra assignment – no credit!

https://canvas.harvard.edu/courses/166148/quizzes/460223

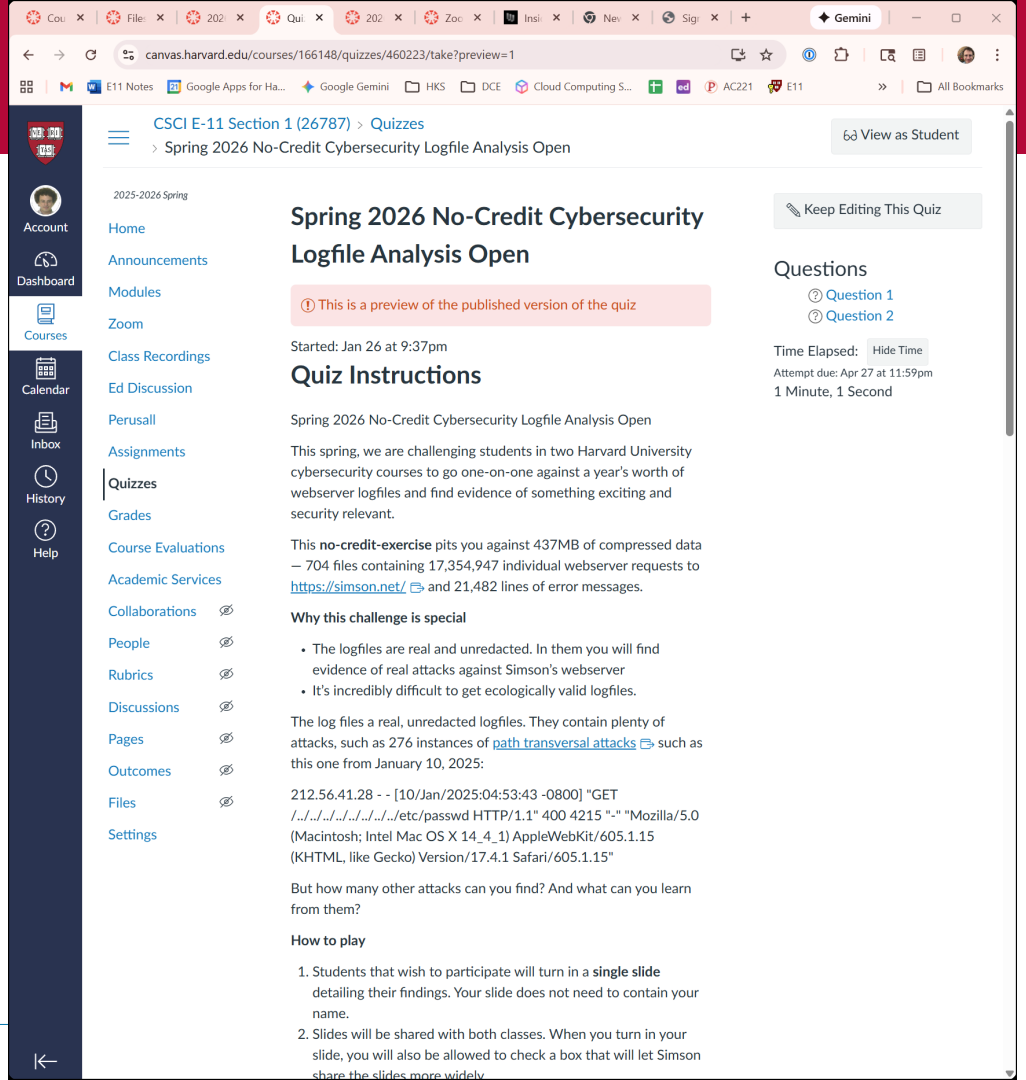

For Next Week

How to read a research paper

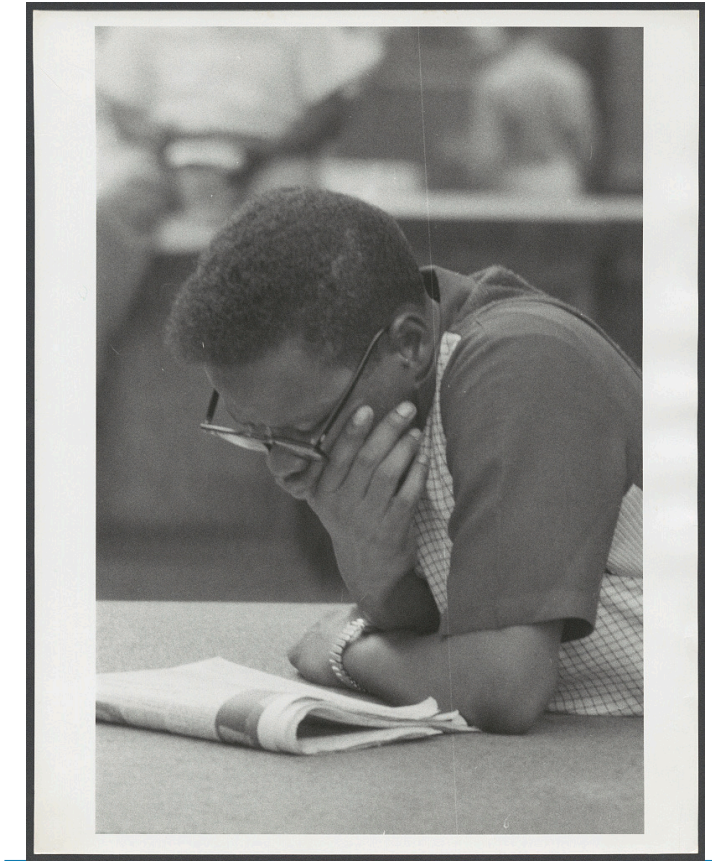

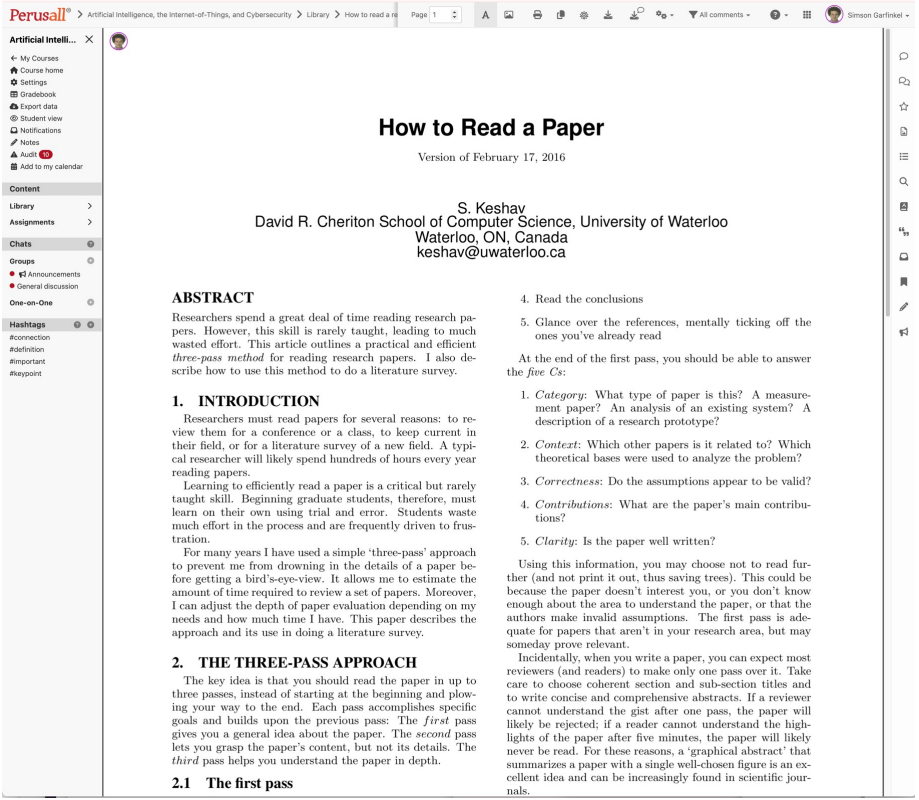

Perusall® > Artificial Intelligence, the Internet-of-Things, and Cybersecurity > Library > How to read a re

Page 1

All comments

Simson Garfinkel

Artificial Intelli...

My Courses
Course home
Settings
Gradebook
Export data
Student view
Notifications
Notes
Audit 10
Add to my calendar

Content
Library
Assignments

Chats

Groups
Announcements
General discussion

One-on-One

Hashtags
#connection
#definition
#important
#keypoint

How to Read a Paper

Version of February 17, 2016

S. Keshav
David R. Cheriton School of Computer Science, University of Waterloo
Waterloo, ON, Canada
keshav@uwaterloo.ca

ABSTRACT

Researchers spend a great deal of time reading research papers. However, this skill is rarely taught, leading to much wasted effort. This article outlines a practical and efficient *three-pass method* for reading research papers. I also describe how to use this method to do a literature survey.

1. INTRODUCTION

Researchers must read papers for several reasons: to review them for a conference or a class, to keep current in their field, or for a literature survey of a new field. A typical researcher will likely spend hundreds of hours every year reading papers.

Learning to efficiently read a paper is a critical but rarely taught skill. Beginning graduate students, therefore, must learn on their own using trial and error. Students waste much effort in the process and are frequently driven to frustration.

For many years I have used a simple 'three-pass' approach to prevent me from drowning in the details of a paper before getting a bird's-eye-view. It allows me to estimate the amount of time required to review a set of papers. Moreover, I can adjust the depth of paper evaluation depending on my needs and how much time I have. This paper describes the approach and its use in doing a literature survey.

2. THE THREE-PASS APPROACH

The key idea is that you should read the paper in up to three passes, instead of starting at the beginning and plowing your way to the end. Each pass accomplishes specific goals and builds upon the previous pass: The *first* pass gives you a general idea about the paper. The *second* pass lets you grasp the paper's content, but not its details. The *third* pass helps you understand the paper in depth.

2.1 The first pass

4. Read the conclusions
5. Glance over the references, mentally ticking off the ones you've already read

At the end of the first pass, you should be able to answer the *five Cs*:

1. *Category*: What type of paper is this? A measurement paper? An analysis of an existing system? A description of a research prototype?
2. *Context*: Which other papers is it related to? Which theoretical bases were used to analyze the problem?
3. *Correctness*: Do the assumptions appear to be valid?
4. *Contributions*: What are the paper's main contributions?
5. *Clarity*: Is the paper well written?

Using this information, you may choose not to read further (and not print it out, thus saving trees). This could be because the paper doesn't interest you, or you don't know enough about the area to understand the paper, or that the authors make invalid assumptions. The first pass is adequate for papers that aren't in your research area, but may someday prove relevant.

Incidentally, when you write a paper, you can expect most reviewers (and readers) to make only one pass over it. Take care to choose coherent section and sub-section titles and to write concise and comprehensive abstracts. If a reviewer cannot understand the gist after one pass, the paper will likely be rejected; if a reader cannot understand the highlights of the paper after five minutes, the paper will likely never be read. For these reasons, a 'graphical abstract' that summarizes a paper with a single well-chosen figure is an excellent idea and can be increasingly found in scientific journals.

Optional reading on Perusall: *How to Read a Paper*

This Week's Reading: Toward Total-System Trustworthiness, Peter G. Neumann

Part 1 of 1 1

viewpoints

DOI:10.1145/3532631

Peter G. Neumann

Inside Risks

Toward Total-System Trustworthiness

Considering how to achieve the long-term goal to systemically reduce risks.

COMMUNICATIONS' INSIDE RISKS columns have long stressed the importance of total-system awareness of riskful situations, some of which may be very difficult to identify in advance. Specifically, the desired properties of the total system should be specified as requirements. Those desired properties are called emergent properties, because they often cannot be derived solely from lower-layer component properties, and appear only with respect to the total system. Unfortunately, additional behavior of the total system may arise—which either defeats the ability to satisfy the desired properties, or demonstrates that the set of required properties was improperly specified.

In this column, I consider some cases in which total-system analysis is of vital importance, but generally very difficult to achieve with adequate assurance. Relevant failures may result from one event or even a combination of problems in hardware, software, networks, operational environments, and of course actions by administrators, users, and misusers. All of the interactions among these entities need to be considered, evaluated, and if potentially deleterious, controlled by whatever means available. The problem to be confronted here is trying to analyze an entire system as a composition of its components, rather than just considering its components individually. In many cases, system failures tend to arise within the interactions and interdependencies among these components, depending on whether a system was designed modularly to minimize disruptive dependencies, with each module carefully specified.

Addressing this problem is a daunting endeavor, even for seasoned developers of critical systems. Whether explicitly defined or implicit, total-system requirements may be highly interdisciplinary, including stringent life-critical requirements for human safety; system survivability with reliability, robustness, resilience, recovery, and fault tolerance; many aspects of security and integrity; guaranteed real-time performance; forensics-worthy accountability, high-integrity evidence, and sound real-time and retrospective analysis; defenses against a wide ranges of physical, electronic, and other adversities; and coverage of numerous potential risks. Ideally, requirements should be very carefully specified at various architectural layers, preferably formally as much as possible in newly developed systems, and especially in particularly vulnerable components. Although this is usually not applicable to legacy systems, it is stated here as a farsighted goal for future developments.

Geoff Cohen just now

How long term is "long-term"? What's the time-line?

Enter your comment or question and press Enter. Mention someone by typing @. Add hashtags by typing #. Quote someone in your reply by selecting text in one of the comments above.

Post comment

Notify me of updates to this conversation

Locate in content

Last viewed thread

Next thread

Next unread thread

Optional Reading: ACM Risks Forum Quarterly Summary, October 2025

ACM Risks Forum Quarterly Summary

Peter G. Neumann
SRI Computer Science Lab
333 Ravenswood Avenue, Menlo Park CA 94025-3493
neumann@CSL.sri.com
Quarter ending 31 August 2025
https://www.CSL.sri.com/users/neumann
DOI: 10.1145/3772008.3772009
https://doi.org/10.1145/3772008.3772009

ABSTRACT

This is an annotated summary of the main items contributed to the ACM Risks Forum in the most recent quarter, orgnanized categorically. References (R i j) to the online Risks Forum denote RISKS vol i number j. Cited RISKS items generally identify contributors and sources, together with URLs. Official RISKS archives are available at www.risks.org, with nice html formatting and search engine courtesy of Lindsay Marshall at Newcastle: `http://catless.ncl.ac.uk/Risks/i.j.html` (also ftp://www.sri.com/risks). *CACM* Inside Risks: `https://www.csl.sri.com/users/neumann/insiderisks.html`

MINI-EDITORIAL (PGN)

Lots more ugly business-as-usual items. If only we had more trustworthy computer systems and better programmers and fewer would be AI 'solutions where they are often the wrong solution to the wrong problem. But many things we depend on seem broken. Please hang in there and contribute rekevant items to RISKS. PGN

TOTAL-SYSTEM TRUSTWORTHINESS

System Trustworthiness encompasses requirements such as system safety, security, reliability, survivability, guaranteed real-time requirements, and assurance that these requirements are met. The items in this section typically require system requirements from multiple disciplines, and thus deserve more holistic attention.

- A brazen attack on air safety is underway. Here's what's at stake. (R 34 75)
- Alaska Airlines Grounds All Flights for Three Hours Due to IT Outage (R 34 73)
- North American Airlines Targeted by Cyberattacks (R 34 70)
- U.S. air-traffic control still runs on Windows 95 and floppy disks (R 34 68)
- Why Can't Americans Buy the Best Electric Vehicle? (R 34 70)
- Tesla blows past stopped school bus and hits kid-sized dummies in n Full Self-Driving tests (R 34 68)
- Tesla Found Partly to Blame for Fatal Autopilot Crash (R 34 75) Tesla Testing if Its Robotaxis Can Be Hacked Remotely (R 34 74); Telsa's robotaxi rollout reported to be a mess (R 34 69); Elon Musk's Tesla hits a speed-bump in California (R 34 76)
- Finger Grease Mitigation for Tesla PIN Pad (R 34 68)
- Tesla Wall Connector Charger Hacked Through Charging Port (R 34 69)
- Tesla Robotaxi Involved in 1st Official Accident (R 34 70)
- Tesla robotaxi incidents spark confusion and concerns in Austin, TX (R 34 70)
- Chinese-made self-driving trucks: Even after it hit a motorcycle and caused accident, it is still running? (R 34 75)
- Mercedes-Benz will let you use an in-car camera in Microsoft Teams while driving (R 34 72)
- Growing Challenge of Radar Interference in Autonomous Vehicles (R 34 72)
- Ford Breaks Annual Record for Safety Recalls Within First Six Months of Year (R 34 70)
- Driving assistance systems could backfire: Some warning alerts can lead to more hazardous driving (R 34 72)
- Railroad industry first warned about this nasty vulnerability in xsx2005. (R 34 72,75)
- Cargo Ship That Caught Fire Carrying Electric Vehicles Sinks in the Pacific (R 34 69)
- Financial institutions should prepare for subsea cable sabotage (R 34 72)
- How nuclear war could start: To understand how it could all go wrong, look at how it almost did. (R 34 68)
- They Asked ChatGPT Questions. The Answers Sent Them Spiraling. (R 34 68)
- New Research Finds That ChatGPT Secretly Has a Deep Anti-Human Bias (R 34 75)
- News Sites Are Getting Crushed by Google's New AI Tools (R 34 68,70)
- Can AI safeguard us against AI? One of its Canadian pioneers thinks so (R 34 68)
- ChatGPT goes down – and fake jobs grind to a halt worldwide (R 34 68)

Start the Labs!

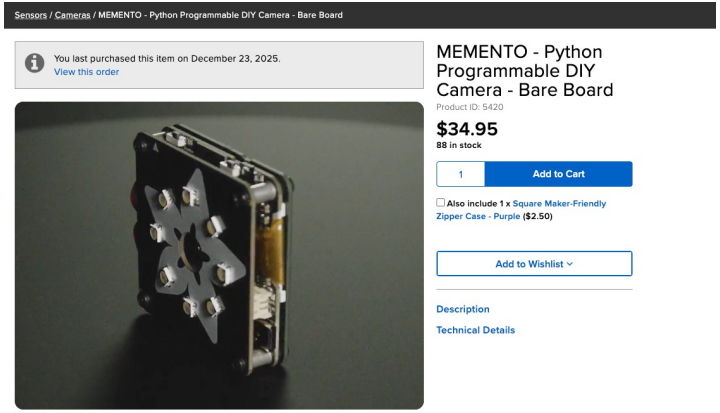

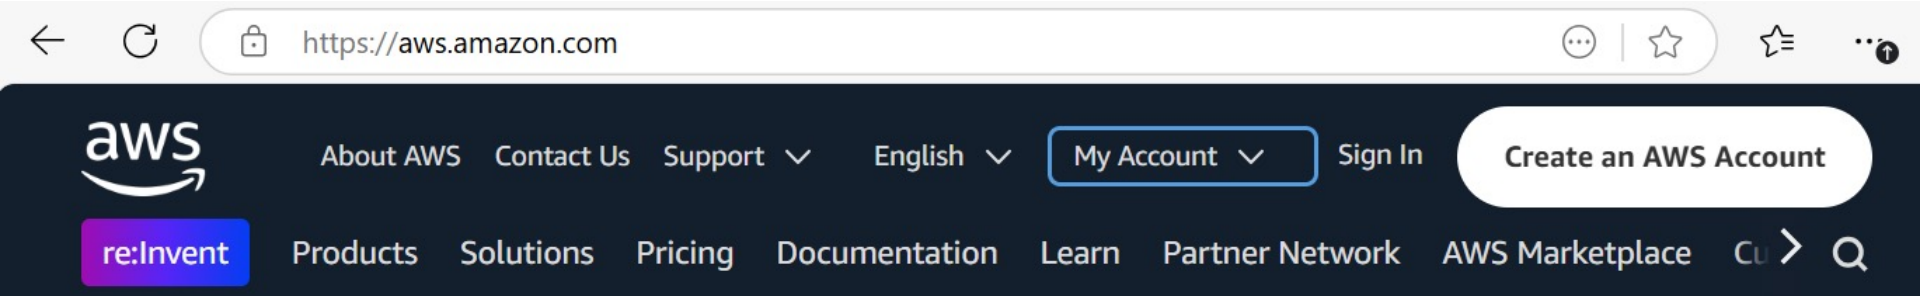

Q&A and Conclusion

Thanks for coming to the first class

For next week:

- Fill out the Welcome Survey
- Fill out the SA-6 Survey
- Introduce yourself on Ed Board
- Participate in the Ed Board discussions
- Do the reading on Perusall and comment on the paper
- Do the Quiz for the first paper
- (Do the optional reading)
- Complete Lab #0
- Start Lab #1 and get as far as you can (Due Feb. 9). Do not put it off!