NC STATE UNIVERSITY

CSC-591/791

LLMs in Security

Reading papers

Alexandros Kapravelos
akaprav@ncsu.edu

How to pick which papers to read?

How to build your bibliography

- Identify the top venues of security research
  - csrankings.org
  - Google Scholar
  - Guofei Gu list
- Identify top papers/people in the field
  - https://www.sec.cs.tu-bs.de/~konrieck/topnotch/
  - http://s3.eurecom.fr/~balzarot/notes/top4_v3/
  - https://nebelwelt.net/pubstats/
- Keep track of top venues/papers/people over time
  - Use github to collaboratively build a more comprehensive bibliography for LLM security
  - Markdown + bibtex
  - Learn how to maintain a knowledgebase of the papers you read
  - How can you have a system so that you never have to read a paper again?

LLM papers

https://arxiv.org/

https://paperswithcode.com/

https://github.com/Hannibal046/Awesome-LLM

Reading a research paper 1/2

- Why are you reading this paper?
  - Your time is limited and you cannot read all literature
  - Decide if the paper is useful
    - Now -> read it
    - Later -> keep it in a “read later” folder
    - Or skip it!
- **Reading for breadth**
  - First skim the paper
  - Intro, section headings, tables & figures, conclusions
  - Is it credible? Top-tier conference, well-known authors, outdated?
  - Skim bibliography, is it complete?
  - How useful is it?
  - Decide whether to go on
  - This process will allow you to follow the paper better

Reading a research paper 2/2

- **Reading for depth** -> Challenge what you read
  - How did they do it?
  - Challenge their arguments
  - Examine assumptions
  - Examine methods
  - Examine statistics/results
  - Examine reasoning and conclusions
- Once you understand the paper, ask yourself how you can apply their approach to your own work

- Take notes as you read
- Highlight major points
- Write a summary
- **If you go back to the paper a year later, can you quickly figure out the major points without reading it again?**

The conference review process

- Paper is submitted to conference
- The technical program chair(s) assign the paper to two or more technical program committee members
- The TPC members provide their reviews
  - Online discussion
  - TPC meeting
- In some conferences there are multiple rounds
  - Round 1 -> two reviews, if all negative early reject
  - If not rejected you might be asked to provide a rebuttal to the first reviews
  - Round 2 -> additional reviews
- Decision (varies a bit, but the basic gist is the following)
  - Bottom third -> rejected
  - Top third -> accepted
  - Papers in the middle -> discussed/major revision

Your reviews for this course 1/2

- You need to start observing the **writing** of the papers and comment on the good parts and what can be improved
- This will allow you to develop a better writing style yourself
- There is going to be a special section in your review system to comment on writing
- Observe also the structure of the papers that you understood well and the structure of the papers that were more difficult to comprehend

Your reviews for this course 2/2

- You need to comment also the technical parts
- Feedback to the authors on how to improve the paper or otherwise proceed with their work

Review format for our class:

- Rating
- Reviewer confidence
- Summary
- Strengths
- Weaknesses
- Detailed comments for authors
- Writing/paper structure comments
- What did you learn from this paper?
- What questions do you have about the paper or the area?