



CNS Research Review Agenda, Thursday, May 8, 2025

8:30 a.m. Breakfast [Computer Science and Engineering building at UC San Diego, Room 1242]

9:00 a.m. Welcome and Round Table Introductions [CSE Room 1242]

Stefan Savage and George Porter, CNS Co-Directors and CSE Professors

9:30 a.m. TBD - Session Chair

• **Characterizing the MrDeepFakes Sexual Deepfake Marketplace (30 min.)**

Presenter: Deepak Kumar, CNS/CSE Faculty

• **Practical Support for Integrity Validation of Criminal Legal Process (15 min.)**

Presenter: Alisha Ukani, CNS/CSE Graduate Student (Alex Snoeren)

• **Count of Monte Crypto (15 min.)**

Presenter: Alex Liu, CNS/CSE Graduate Student (Stefan Savage and Geoff Voelker)

10:30 a.m. Break and Informal Interaction (15 minutes)

10:45 a.m. TBD - Session Chair

• **Inferring Hospital Ransomware Outages (15 min.)**

Presenter: Sumanth Rao, CNS/CSE Graduate Student (Stefan Savage and Geoff Voelker)

• **Understanding the Efficacy of Phishing Training in Practice (30 min.)**

Presenter: Elisa Luo, CNS/CSE Graduate Student (Stefan Savage and Geoff Voelker)

• **Setting Up Emergency Hospital Networks (15 min.)**


Presenter: Almog Bar-Yossef, CSE Student, and Kartikeyan (Kartik) Subramanyam, Network Systems Engineer, UC San Diego Health

11:45 a.m. Discussion Group Introduction - Alisha Ukani

Group Photograph [Bear Sculpture outside of CSE]

12:15 p.m. Discussion Topics and Informal Interaction (30 minutes) Small discussion groups where students and industry members can chat about open research questions on different topics.

12:45 p.m. Lunch [CSE Room 1202]



CNS Research Review Agenda (Page 2), Thursday, May 8, 2025

2:00 p.m. TBD - Session Chair

- **Alarming Alarms: Analyzing the Security of Dealer-Installed Car Alarm Systems (30 min.)**

Presenter: Yibo Wei, CNS/CSE Graduate Student (Aaron Schulman)

- **Wide-area Carbon Load of Network Transfers (15 min.)**

Presenter: Amanda Tomlinson

- **Breaking Enterprise Network Authentication (30 min.)**

Presenter: Miro Haller, CNS/CSE Graduate Student (Nadia Heninger)

3:15 p.m. Break and Informal Interaction (30 minutes)

3:45 p.m. TBD - Session Chair

- **Automated Verification 101: How Verification Impacts Software Engineering (30 min.)**

Presenter: Eric Mugnier, CNS/CSE Graduate Student (YY Zhou)

- **Sublet Your Subnet: Inferring IP Leasing in the Wild (15 min.)**


Presenter: Ben Du, CNS/CSE Graduate Student (Alex C. Snoeren, kc claffy)

- **Abacus: Type Enforced Peripheral Management (30)**

Presenter: Tyler Potyandy, CNS/CSE Graduate Student (Pat Pannuto)

5:00 p.m. Lightning Talks – Graduate Students and Industry Representatives (30 min.)

6:00 p.m. Graduate Student Poster Session and Dinner Reception [Great Hall @ UCSD]



Day Two / CNS Research Review Agenda, Friday, May 9, 2025

9:00 a.m. Breakfast [Computer Science and Engineering building at UC San Diego, Room 1242]

9:30 a.m. TBD - Session Chair

- **Backhauling Data via Aerial Infrastructure (15 min.)**

Presenter: Alex Yen, CNS/CSE Graduate Student (Pat Pannuto)

- **Satellite Monitoring (30 min.)**

Presenter: Wenyi (Morty) Zhang, CNS/CSE Graduate Student (Aaron Schulman)

- **Fun-tuning: Characterizing the Vulnerability of Proprietary LLMs to Optimization-based Prompt Injection Attacks via the Fine-Tuning Interface (15 min.)**

Presenter: Andrey Labunets, CNS/CSE Graduate Student (Earlence Fernandes)

10:30 a.m. Break and Informal Interaction (15 minutes)

10:45 a.m. TBD - Session Chair

- **An Empirical Analysis of the Use of and Reporting of National Security Letters (30)**

Presenter: TBD

- **Talk TBD (30)**

Presenter: TBD

11:45 a.m. Open Floor [Feedback from company representatives]

12:00 p.m. Lunch and CNS Research Review Conclusion [CSE Room 1202]

-CNS Research Review Conclusion-

-See Abstracts on Next Page-

Abstracts in order of appearance:

Practical Support for Integrity Validation of Criminal Legal Process - Alisha Ukani

In the US, criminal legal process (e.g., search warrants, subpoenas, etc.) is the key mechanism by which law enforcement entities and courts compel third parties (e.g., such as Google or Meta) to produce evidence or take actions in support of a criminal investigation. However, such orders are simply documents and there is no inherent mechanism to guarantee their integrity. In recent years there have been increasing reports of forged legal process used to advance a variety of criminal purposes, with online sellers now offering such capabilities for a fee.

Our work focuses on designing a practical system to help third parties identify when the legal orders they receive are the result of forgery or tampering. Based on a series of conversations with judges, court clerks, law enforcement and the legal compliance components of large online service providers, we have identified a set of security and functionality goals that capture both the integrity needs of this ecosystem and the reality of how legal process is created, modified and served. Based on these, we have designed a system, to be deployed at both issuers (e.g., courthouses) and providers (e.g., Google) to help authenticate and validate the integrity of such documents after they are served.

Count of Monte Crypto - Alex Liu

Between 2021 and 2023, crypto assets valued at over \$US2.6 billion were stolen via attacks on “bridges”—decentralized services designed to allow inter-blockchain exchange. While the individual exploits in each attack vary, a single design flaw underlies them all: the lack of end-to-end value accounting in cross-chain transactions. In this paper, we empirically analyze over 10M million transactions used by key bridges during this period. We show that a simple invariant that balances cross-chain inflows and outflows is compatible with legitimate use, yet precisely identifies every known attack (and several likely attacks) in this data. Further, we show that this approach is not only sufficient for post-hoc audits, but can be implemented in-line in existing bridge designs to provide generic protection against a broad array of bridge vulnerabilities.

Inferring Hospital Ransomware Outages - Sumanth Rao

Ransomware attacks affecting hospitals pose a substantial risk to patient safety, as well as the financial and operational viability of the national healthcare infrastructure. Current regulatory framework and status quo behavior result in delayed reporting of ransomware attacks by the hospitals themselves. This causes a detrimental effect on the ability to rapidly deploy resources to affected institutions, and in preparing unaffected hospitals within the “cyber blast radius” for the increased patient diversion and workflow burdens. This talk presents an ongoing effort to develop a measurement tool which can be used to infer ransomware-induced hospital network outages. Using this, we characterize ransomware attacks in the past and other outages affecting hospital organizations.

Understanding the Efficacy of Phishing Training in Practice - Elisa Luo

Our work empirically evaluates the efficacy of two ubiquitous forms of enterprise security training: annual cybersecurity awareness training and embedded anti-phishing training exercises. Specifically, our work analyzes the results of an 8-month randomized controlled experiment involving ten simulated phishing campaigns sent to over 19,500 employees at a large healthcare organization. Our results suggest that these efforts offer limited value. First, we find no significant relationship between whether users have recently completed cybersecurity awareness training and their likelihood of failing a phishing simulation. Second, when evaluating recipients of embedded phishing training, we find that the absolute difference in failure rates between trained and untrained users is extremely low across a variety of training content. Third, we observe that most users spend minimal time interacting with embedded phishing training material in-the-wild; and that for specific types of training content, users who receive and complete more instances of the training can have an increased likelihood of failing subsequent phishing simulations. Taken together, our results suggest that anti-phishing training programs, in their current and commonly deployed forms, are unlikely to offer significant practical value in reducing phishing risks.

Alarming Alarms: Analyzing the Security of Dealer-Installed Car Alarm Systems - Yibo Wei

Our analysis of three popular aftermarket alarm ecosystems revealed a critical vulnerability in one major system that allows remote attackers to control vehicle functions including the immobilizer, doors, and CAN bus. With an estimated 1.68 million vulnerable alarms deployed, we recommend combining security design elements from multiple systems to achieve defense in depth.

Breaking Enterprise Network Authentication - Miro Haller

The RADIUS protocol is the de facto standard lightweight protocol for authentication, authorization, and accounting for networked devices. It is used to support remote access for diverse use cases including network routers, industrial control systems, VPNs, enterprise Wi-Fi including the Eduroam network, Linux Pluggable Authentication Modules, and mobile roaming and Wi-Fi offload. This talk presents the Blast-RADIUS vulnerability which allows a machine-in-the-middle attacker on RADIUS to authenticate themselves to a device as an arbitrary user with privileges of their choosing. We discuss the proof of concept applications of our attack against popular RADIUS implementations, and the large-scale disclosure process and mitigation efforts in collaboration with CERT and IETF.

Sublet Your Subnet: Inferring IP Leasing in the Wild - Ben Du

In this talk, we analyze the IPv4 leasing ecosystem by designing a methodology to infer leased address space globally and study its impact on routing and hosting security. We infer that 4.1% of all advertised IPv4 prefixes (0.9% of routed v4 address space) were leased in April 2024. We show that leased address space is five times more likely to be abused compared to non-leased space.

Automated Verification 101: How Verification Impacts Software Engineering - Eric Mugnier

Verification has never been as popular as today. While the excitement is palpable, feedback on the practical use of verifiers remains scarce. We lack a clear understanding of how these tools are employed, the challenges users encounter, and which types of projects are well-suited for verification. Although some insights can be gleaned from blog posts, and talks, these are individual anecdotes rather than a comprehensive view of the field.

In our work, we address these gaps by interviewing several industrial users of verification tools, focusing on how the use of an automated verifier—specifically Dafny—impacted the software development process. Using a method inspired by the sociology technique known as grounded theory, which develops theories and findings directly from data, we aim to help the community better understand and effectively utilize these tools. Additionally, we interviewed the developers of Dafny to gain their perspective on the tool's use, lessons learned, and potential future directions. Through this work, we aim to provide a comprehensive view of automated verifiers, including insights on suitable projects, best practices and pitfalls, and motivations for future research.

Fun-tuning: Characterizing the Vulnerability of Proprietary LLMs to Optimization-based Prompt Injection Attacks via the Fine-Tuning Interface - Andrey Labunets

We surface a new threat to closed-weight Large Language Models (LLMs) that enables an attacker to compute optimization-based prompt injections. Specifically, we characterize how an attacker can leverage the loss-like information returned from the remote fine-tuning interface to guide the search for adversarial prompts. The fine-tuning interface is hosted by an LLM vendor and allows developers to fine-tune LLMs for their tasks, thus providing utility, but also exposes enough information for an attacker to compute adversarial prompts. Through an experimental analysis, we characterize the loss-like values returned by the Gemini fine-tuning API and demonstrate that they provide a useful signal for discrete optimization of adversarial prompts using a greedy search algorithm. Using the PurpleLlama prompt injection benchmark, we demonstrate attack success rates between 65% and 82% on Google's Gemini family of LLMs. These attacks exploit the classic utility-security tradeoff - the fine-tuning interface provides a useful feature for developers but also exposes the LLMs to powerful attacks.