Northeastern University
EECE 5698 Fall 2025
Design Project:
Protecting the Security of Voice-Controlled Systems from Laser Injection Attacks

Instructor: Prof. Kevin Fu
TAs: Hui Zhuang, Nuntipat Narkthong
Last updated: 10/20/2025, 3 PM

Project Title Deadline: **During lecture on Thu, Oct 9**
Oral Proposal Deadline: **During lecture on Mon, Oct 20**
Written Proposal Deadline: **Thu, Oct 30 11:45 AM**
Milestone Presentation Deadline: **During lecture on Thu, Nov 20**
Project Demo: **In open lab hours Nov 24 - 28.** Timeslot will be posted later.
Tournament Deadline: **During the lecture on Thu, Dec 4**
Final Report Deadline: **Mon, Dec 8 11:45 AM**

0. Overview

The team project investigates in depth the consequences of laser-based speech audio injection attacks against actual voice-controlled systems as well as the possible defenses against such attacks. The project consists of three implementation components:

- $\mathcal{C}1$: Building voiced-controlled systems that can be used in real-world applications.
- $\mathcal{C}2$: Developing laser injection methods for attacking these voice-controlled systems.
- $\mathcal{C}3$: Developing defense mechanisms that protect your systems against signal injection attacks.

In practice, of course, one would want a Security Development Lifecycle where defense is designed into a product rather than bolted on after the fact. In addition to these three components, you will also provide a written, thoughtful threat model plus a performance and security evaluation mechanism for your specific systems so that other student teams can cross-evaluate your system designs. Students in class will form **teams of three (3)** and will participate in a final tournament that reveals each team's project achievements by gaming with each other as defenders and attackers. You are free to choose your teammates, but the staff reserves the right to adjust team membership to ensure good conditions for the tournament. In the end, students' project grades will be based on three aspects that correspond to $\mathcal{C}1$, $\mathcal{C}2$, and

C3, namely (1) their voice-controlled system's performance in the proposed application (without attacks), (2) students' ability to attack systems built by others, (3) their system's robustness against attacks launched by other students.

Budget

Our lab has spare microcontrollers (such as raspberry pi) and potentially other parts that you may need. Please let TAs know your needs on Piazza by the proposal deadline. Remember that first come, first served. We will allow students to spend up to $100 of their own additional funds. So be frugal, and be creative. We will expect teams to provide a spreadsheet of any additional purchases to justify that costs were kept under $100. Finally, if you think there are needs of devices/components that are shared among all teams and may be potentially purchased by instructors in bulk, please post it on Piazza so that all students and instructors know them.

Note: supply chains are taking longer than usual for purchasing, so it's crucial to purchase any exotic components immediately. In the past, some purchases have taken weeks and months to arrive on campus.

1. Project Proposal

The project proposal mostly focuses on ***C1***. Teams will give an oral proposal in lecture (see header) and submit a formal written proposal with detailed plans.

1.1 Oral Proposal

During a lecture, teams will propose the **title** of their projects to the instructors. We will later dedicate one lecture hour to a sequence of private team presentations. Each team will privately brief the instructors on their planned project proposal. A team may present slides, use the whiteboard, and speak for **up to 5 minutes**, followed by Q/A and feedback from the instructional staff. The oral proposal should include a project description of the problem being solved with the voice assistant (see the application section of the written proposal), a list of team members, a team name, and a RACI chart showing each team member committed as responsible for at least one item and accountable for at least one item. We will look for balance in workload among roles, and no role should be heavily weighted chronologically at one part of the semester; each student should have an active role for the entire semester.

Before the oral proposal deadline, students should also read the laser injection paper Light Commands and come up with preliminary ideas for attack and defense approaches (***C2***, and ***C3***) to discuss with instructors during the lecture for oral proposals. **In the final tournament, each team needs to demonstrate two different defense methods: one software-centered defense, and hardware-involved defense that improves the signal conditioning chain.**

After the oral project proposal meeting and incorporating instructors' feedback, students should prepare the formal written proposal and start the actual implementation of ***C1***, ***C2***, and ***C3***.

1.2 Written Proposal

Your formal written proposal should contain the following information. It should be no more than five (5) pages, double spaced inclusive of text, illustrations, and references.

Application

Choose an existing application (or a specific system) that can utilize the proposed voice control functionalities to enhance its efficiency and/or usability. Some examples include controlling medical devices, automobiles, stock trades, home security, home automation, and drones with voice. You should also provide technical reasonings backed up by references to explain why voice controls can enhance this existing application. In your implementation, you are welcome to be creative with toys or partial simulations as there is no budget to buy a car or nuclear powered submarine yet. For the detailed requirement of your system application, please see Section 6.

Threat Modeling

Provide a thoughtful threat model for your system that centers on sensor attacks, with particular emphasis on laser injection attacks. You should consider both lower-level physical attack scenarios and higher-level application layer consequences.

Interface & Testing Procedure

Provide a detailed specification of the voice control interface you plan to add to the existing application including a list of usable voice commands that legitimate users and potential adversaries can use to test your system. Based on the interface, you should further propose a testing procedure that allows objective evaluation of your voice-controlled system. Your testing procedure should capture the notions of performance (without attacks) and security (under attacks) of your designed systems. Furthermore, the testing procedure should be centered on objective metrics and provide more than one level of granularity for performance analysis (e.g., the rate of successfully recognized voice commands alone is insufficient).

2. Project Milestone Presentation

Students should at least have finished ***C1*** at this milestone. The project milestone presentation will be held during lecture, where students present the designs and implementations of their voice controlled systems to instructors and other teams of students. The designs and outcomes of $\boldsymbol{C2}$ and $\boldsymbol{C3}$, assuming already obtained at this point, will not be presented yet. More specifically, we expect the following deliverables in the presentation:

- Contents included in the project proposal excluding the threat model which is also considered as classified information.
- A live or video demo of the working system.

Tournament Announcement

Instructors will announce the detailed rules, procedures, and metrics of the upcoming tournament after the milestone presentations. Be prepared to make further tweaks to your systems so that your systems and interfaces can be adapted to the fast-paced tournament.

3. Project Demo

Teams will present their project designs and give live demos during demo week. Each team has **30 min** including any setup time and should demo the following components:

- How normal human speech and DTMF tones interact with their systems, including the three functional commands required by the tournament.
- How audio injection attacks interact with their systems when there is no defense mechanisms applied.
- How different defense mechanisms applied to the system defend against the audio injection attacks.

Besides the live demos, we also encourage teams to make a few slides that clearly show any unique elements to their system and overviews of their attack and defense methods without explaining the implementation details. All other teams should come to these demo sessions to learn the defense mechanisms so they can start to devise their attacks for the tournament.

4. Tournament

The teams will compete in a two-round tournament with 3 games in total. In each game, one team acts as the system designer and the other one acts as the attacker and they will then swap their roles. Winners will get prizes. We will design the tournament procedure to keep each game within **35 min including any setup time** (team A and B act as the adversary and designer/defender in the first 17 min and then swap around in the second half). The achievement of $\boldsymbol{C1}$ and $\boldsymbol{C2}$ will be evaluated when acting as the designer; $\boldsymbol{C3}$ will be evaluated when performing the attacker role. There will be 3 games in total.

5. Project Paper

Your final project paper should explain the details of system design and implementation, provide self-evaluations of your system's performance and security by using the testing procedure you proposed before, and analyze your system's performance in the tournament and possible future improvements.

6. System Requirements

- You must use ADMP401 as the microphone. We will provide one module to each team for free.
- It should allow different people's voices to interact with it. That is, the voice recognition model needs to support speaker-independent recognition.
- Your voice control assistant may leverage open source software. Your system should minimally recognize three different, interesting voice commands (e.g., "start car" or "stop car" or "accelerate"). There should also be at least one voice command for authentication of a command (e.g., "My name is Kevin and my voice is my passport. Verify me."
- Your system should accept voice commands with DTMF tones to represent digits 0-9 for an interesting function (e.g., "Set speed. 9#" or "Deposit into a bank account. 1234#" etc.).
- To make the game more interesting, you are not allowed to use cryptography. This forces your system to have certain intrinsic insecurities for other teams to exploit. However, you can try to deploy compensating controls to make up for the horrible decision of not to allow cryptography.
- The system should have two defense methods as mentioned in Section 1.1. Further, the defenses should be controlled by on/off switches, which means they can be enabled and disabled during the tournament dynamically.
- Your attacks should be non-destructive, i.e., cannot physically break your opponent's systems.

7. Grading Rubrics

All members in a team will receive the same grade. However, all students are subject to the Academic Integrity Policy at Northeastern, and it would be an policy violation to misrepresent your role or contributions which ought to be fairly distributed amongst the team. To receive a grade, all borrowed equipment and the $100 budgeted purchases must be returned or replaced/reimbursed. No grade shall be issued to a student until such equipment is returned or replaced/reimbursed.

Proposal Grading

- 40% Applications
- 30% Quality and depth of threat model
- 30% Self testing procedure

Milestone Presentation + Project Demo Grading

- 35% Conveyance of meaningful content with intellectual depth (a compelling problem, approach, and demonstration)
- 15% Lively slide and/or whiteboard format/structure catered to your topic (e.g., not a cookie-cutter list of endless bullets)
- 15% Engagement of audience (whether your fellow students learn something new)
- 10% Structured presentation (intro, structured body, conclusion)
- 10% Delivery of presentation (speaking skills)
- 10% Organized and exacting use of time (not ending early, not ending late)
- 5% Overall impression

Project Paper

- 40% Content (adherence to proposed topic, sufficient development of thesis)
- 25% Mechanics of writing (sentence structure; clarity; cohesive paragraphs; accuracy in usage, grammar, and spelling)
- 15% Originality, creativity of the defense and attack approach
- 10% Related work (formal, alphabetized citations as found in CS research papers)
- 5% Content structure (abstract, intro, background, related work, your approach, conclusion)
- 5% Overall impression

Addendum 1

Tournament Procedure and Rubric

- The microphone and laser are 20 cm away. We will provide three DC power supplies for each attacker-defender group.
- Each team picks three functional commands, with one command containing DTMF tones. In the demo week, each team’s demo needs to include the three voice commands that will be used in the tournament. If you use another authentication command as a defense, that authentication command is considered part of the three functional commands.
- All defenses need to be able to be turned off. Otherwise, the attackers win automatically. Other teams are invited to come to the demo evaluations so they can start to devise their attacks for the tournament.
- The two teams flip a coin to decide which team will be the attacker first.

- When there is a tie in the scores, the team that uses a shorter time to finish their attacks wins. If still a tie (neither team finished their attacks), then flip a coin to decide the winner.

For each attacker-defender pair (5 min setup time + 15 min testing)

Part 1: Defenders interact with their own systems when all defenses are turned ON.

- 2 min maximum
- 1 pt for each command to the defender

Part 2: Attackers inject lasers into defenders' systems when all defenses are turned OFF.

- 1 pt for each command to the attacker

Part 3: Attackers inject lasers into defenders' systems when all defenses are turned ON.

- 2 pt for each command to the attacker
- Note: PIN and any other authentications are counted as defenses.

Part 4: Attackers inject other potential physical signals (e.g., EM, ultrasound, etc) into defenders' systems when all defenses are turned ON.

- 2 pt for each command to the attacker

The defender has to carry out part 1 in the 2-min limit, while the attacker can choose to do part 2-4 in any order and with any time allocation strategies within the total 20 min.

Update Log:

Oct 20: Remove the requirement “The functionality of the LED indicators under different audio energy levels” from the Project Demo part.