At Scale and At Risk: Understanding LLM-Generated Vulnerability Reports in Open-Source Ecosystems

Project Description: Open-source projects hosted on GitHub increasingly receive vulnerability reports generated by LLMs. While this lowers the barrier to reporting, it also introduces:

- Hallucinated vulnerabilities
- Fabricated CVE identifiers
- Incorrect severity scores
- Duplicate or low-quality issue submissions
- Maintainer fatigue and triage overload

This project will conduct a large-scale empirical study of LLM-generated vulnerability reports submitted to GitHub repositories. It will:

- Identify patterns of LLM-generated reports
- Measure accuracy against ground-truth sources such as NIST NVD
- Analyze false positive and hallucination rates
- Study maintainer response dynamics
- Evaluate ecosystem impact

The project aims to propose automated validation or filtering mechanisms to support maintainers.

Deliverable: empirical dataset, analysis framework, and a conference paper draft.

Project Type: Research

Internship Batch: Batch 1 or Batch 2

Duties/Activities:

- Collect vulnerability-related issues from GitHub repositories
- Identify LLM-generated patterns (linguistic and structural signals)
- Cross-validate claims with NVD
- Quantify triage burden and false reports
- Propose mitigation strategies
- Draft conference paper

Required Skills:

- Python
- Web scraping / API usage
- Data analysis
- Basic vulnerability management knowledge

Preferred Intern Academic Level: Senior undergraduate or MSc

Learning Opportunities:

- Open-source ecosystem research
- Empirical software security studies
- Large-scale dataset construction
- Research publication process

Expected Team Size: 3 students

Mentors:

- Dr. Yazan Boshmaf (yboshmaf@hbku.edu.qa)
- Dr. Mohannad Alhanahnah (malhanahnah@hbku.edu.qa)