Meet Udeshi

Last updated: Nov. 2026

Ph.D. Candidate, NYU Tandon School of Engineering — Expected graduation: May 2027

m.udeshi@nyu.edu | https://mudeshi.in | github.com/udiboy1209 | Google Scholar | +1-607-216-5524

SUMMARY

Ph.D. candidate in cyber-physical and embedded systems security. I build LLM agents, benchmarks and sandboxes for security analysis (NeurIPS, ICML, AAAI), reverse engineer binaries and study AI-powered malware. Previously three years on Qualcomm's Cloud AI100 ML compiler.

EDUCATION

2022–present **Ph.D., Electrical Engineering**, NYU Tandon. Control/Robotics Research Lab, Center for Cybersecurity. GPA 3.958/4. Advisors: Farshad Khorrami, Ramesh Karri.

2014–2019 **B.Tech. and M.Tech. (Dual Degree), Electrical Engineering**, IIT Bombay. CGPA 8.18/10. Thesis advisor: Virendra Singh.

RESEARCH PROJECTS

LLM Agents for Cybersecurity

Mar 2024–present

Ph.D. research, Center for Cybersecurity, NYU

- Curated NYU CTF Bench (NeurIPS 2024), 200 CTF challenges with an automated evaluation framework for autonomous LLM security agents, and co-authored a lightweight CTF benchmark with LLM-as-a-judge scoring (AAAI 2026).
- Developed EnIGMA (ICML 2025), whose Interactive Agent Tools let LLM agents drive debuggers and other interactive utilities, and D-CIPHER, a planner–executor multi-agent framework; both set the state of the art on NYU CTF Bench.
- Organized the CSAW LLM Attack Challenge (2024–2026), judging agentic solutions from 64 teams.

Cyber-Physical Systems Security

Sep 2025–present

Ph.D. research, Control/Robotics Research Lab, NYU

- Built SaMOSA, an orchestrated QEMU sandbox that observes CPS malware through hardware performance counters, disk, network and syscall activity (EMSOFT 2026), and a sandbox-enabled digital twin for CPS (IOLTS 2026).
- Designed an LLM binary-diffing framework that detects malicious code injected into software updates. Proposed the *Ransomware 3.0* threat model and prototyped *PromptLock*, which ESET found on VirusTotal and initially reported as the first AI-powered ransomware.

Reverse Engineering of Cyber-Physical Systems

Aug 2022–present

Ph.D. research, Control/Robotics Research Lab, NYU

- Designed REMaQE, an `angr` symbolic execution framework that recovers math equations from CPS binaries with 100% accuracy on C and Simulink code (ACM TCPS); with it, found a PID-controller bug in Linux `tools/thermal/tmon` and submitted a Linux kernel patch.
- Designed REMEND, a Transformer neural decompiler with 89.8–92.4% accuracy on C and Fortran binaries, 8% above state-of-the-art neural decompilers on real-world code (ACM TIST).

Tamper-Proof Network Measurements and Hardware Security

Ph.D. research, NYU (IEEE TNSM 2025); M.Tech. thesis, IIT Bombay (2018–2019)

Designed tamper-proof traffic measurement for intrusion detection on a Netronome SmartNIC (P4), out of reach of host-OS rootkits, at 10 Gbps with 99.5% of baseline throughput. Designed cache side-channel attacks exploiting the hardware prefetcher and the re-order buffer.

WORK EXPERIENCE

Siemens Healthineers USA, Princeton, NJ

May 2026–Aug 2026

AI Cyber Intern, AI Factory

ARPA-H UPGRADE: modeled the cyber-clinical impact of unpatched vulnerabilities in hospitals to make patching more efficient. Contributed to LLM-generated synthetic vulnerabilities as RL training data; ran vLLM inference and Slurm-parallel simulations.

Qualcomm R&D, Bangalore, India July 2019–July 2022

Senior Engineer, ML Compiler Team for Cloud AI100 Accelerator

Built multi-core data tiling, memory management, graph scheduling and operator fusion in the AI100 compiler, and kernels for ResNet and YOLO. Worked in PyTorch and TensorFlow op lowering; contributed to *PyTorch Glow*. Deployed a low-power KCF object tracker on AI100.

Amazon India Jul 2017–Aug 2017

Software Development Intern, Transportation Financial Systems

Processed 1 million+ receipts daily (DynamoDB, SQS); automated 30+ AWS resources with CloudFormation.

PUBLICATIONS

Equal contributions are marked with †.

[C1] T. Abramovich, **M. Udeshi**†, M. Shao†, et al. "EnIGMA: Interactive Tools Substantially Assist LM Agents in Finding Security Vulnerabilities." *ICML*, 2025.

[C2] M. Shao†, S. Jancheska†, **M. Udeshi**†, B. Dolan-Gavitt†, et al. "NYU CTF Bench: A Scalable Open-Source Benchmark Dataset for Evaluating LLMs in Offensive Security." *NeurIPS*, 2024.

[C3] M. Shao, N. Rani, K. Milner, H. Xi, **M. Udeshi**, et al. "Towards Effective Offensive Security LLM Agents: Hyperparameter Tuning, LLM as a Judge, and a Lightweight CTF Benchmark." *AAAI Conference on Artificial Intelligence*, 2026.

[C4] **M. Udeshi**, V. S. C. Putrevu, P. Krishnamurthy, R. Karri, and F. Khorrami. "An Orchestrated Sandbox for Side-Channel Analysis of Cyber-Physical Systems Malware." *IEEE Embedded Systems Letters*, presented at *EMSOFT*, 2026.

[C5] **M. Udeshi**†, M. Raz†, P. Krishnamurthy, R. Karri, and F. Khorrami. "Sandbox-Enabled Digital Twin for Cyber-Physical Systems." *IOLTS*, 2026.

[C6] N. K. Boran, P. Diwan, **M. Udeshi**, S. Rathore, and V. Singh. "HIDC: Heterogeneous-ISA Dynamic Core." *HPCC*, 2024.

[J1] **M. Udeshi**, P. Krishnamurthy, R. Karri, and F. Khorrami. "REMEND: Neural Decompilation for Reverse Engineering Math Equations from Binary Executables." *ACM TIST*, 2025.

[J2] **M. Udeshi**, P. Krishnamurthy, R. Karri, and F. Khorrami. "Tamper-Proof Network Traffic Measurements on a NIC for Intrusion Detection." *IEEE TNSM*, 2025.

[J3] **M. Udeshi**, P. Krishnamurthy, H. Pearce, R. Karri, and F. Khorrami. "REMaQE: Reverse Engineering Math Equations from Executables." *ACM TCPS*, 2024.

[J4] N. K. Boran, S. Rathore, **M. Udeshi**, and V. Singh. "Fine-Grained Scheduling in Heterogeneous-ISA Architectures." *IEEE CAL*, 2021.

HONORS AND AWARDS

IEEE CEDA Student Travel Grant (2026); DAC Young Fellowship (2023); Qualcomm RoCStar award (2021); Gold Medal, Indian National Physics Olympiad, top 35 nationwide (2014).

SERVICE AND TEACHING

Peer Review: AAAI, ICLR, ICML and NeurIPS (2026), CPS-IoTSec 2025, and four journals.

2024–2026	Organizer, LLM Attack Challenge, Cybersecurity Awareness Week (CSAW)	NYU
2016–2021	Mentor, Qualcomm Innovation Fellowship; TA for three courses at IIT Bombay.	

SKILLS

Analysis	Symbolic execution (angr), neural decompilation, binary diffing, malware sandboxing
Tools	Ghidra, QEMU, Docker, AWS EC2, Slurm/HPC, LLM agents, vLLM, PyTorch, Fairseq
Programming	Python, Embedded C/C++, Shell/Bash, Assembly, Verilog, P4, LLVM IR
Open source	Linux kernel, Kivy (KivEnt, GSoC 2016), beets, fast_playlist, cpub