Bikrant Bikram Pratap Maurya

bikrant.acc.edu@icloud.com | Portfolio | Google Scholar
LinkedIn: vikrant-vikram-a07418185 | GitHub: vikrant-vikram

Education

M.Tech, Computer Science and Engineering *2024 – 2026*
IIIT-Delhi, India

Publications (Accepted)

DECKER: Domain-Invariant Embeddings for Cross-Keyboard Acoustic Side-Channels | *AsiaCCS 2026* [paper]
Developed a domain-invariant keystroke inference framework using Keyboard Signature Normalization and adversarial disentanglement to extract discriminative features, enabling zero-shot transfer across unseen hardware and environments.

Impact Analysis of Speech Representation Learning Models for Acoustic Side-Channel Attack | *Interspeech 2026*
Introduced the KEYAC dataset and a KAN-based nonlinear adaptation method to establish a new state-of-the-art for keyboard inference attacks against pretrained speech models under severe codec distortion.

GoCoMA: Hyperbolic Multimodal Fusion for LLM Code Attribution | *ICME 2026* [paper]
Proposed a geometry-aware framework that projects code stylometry and binary artifact representations into a hyperbolic Poincaré ball, utilizing geodesic-cosine similarity-based cross-modal attention for robust source attribution across 12 LLM families.

FOCA: Hyperbolic Cross-Attention for Multimodal Malware Classification | *ICASSP 2026* [paper]
Designed a hyperbolic cross-attention mechanism for multimodal malware classification that fuses behavioral and execution traces, achieving state-of-the-art results on cross-environment benchmarks.

CORAL: Socially-Aware Deception Detection via Positive-Unlabeled Learning | *AACL-IJCNLP 2025(workshop)*
Implemented positive-unlabeled learning over conversational social graphs to enable accurate, socially-aware deception detection in sparse-label settings by modeling user interactions and relational context.

Publications (Under Review)

Are Schema-Bound Defenses Sufficient Against Psychology-Informed Jailbreak of LLM Agents? | *Submitted to ACM CCS 2026*
Introduced PsyAgent-Bench to evaluate psychology-informed jailbreaks against multi-agent systems, identifying a Trust-Decay effect and evaluating StoicFirewall as a mitigation layer against semantic steganography.

Surviving the Codec: Continuous Sparse Transduction for Acoustic Side-Channel Attacks over VoIP | *Submitted to Interspeech 2026*
Modeled acoustic side-channel attacks as continuous sparse transduction, proposing EvFlow-CTSSM with Latent Rectified Flow to align non-linearly distorted VoIP embeddings via optimal transport.

Technical Skills & Hobbies

Languages & Frameworks: Python, PyTorch, CUDA, CPP, Javascript,LaTeX
Specialized Concepts: Acoustic Side-Channel Attacks (ASCA), LLM Agent Security & Jailbreaks, Software Forensics, Malware Classification, Semantic Steganography, Adversarial ML
Hobbies: Tennis, strength training, running

Academic Service & Mentorship

Mentorship: Guided student research and undergraduate coursework, providing technical feedback on systems security and rubric-based evaluation for computer science projects.