DANIAL SAMADI VAHDATI

Philadelphia, PA | Open to Relocation
danielsvahdati@gmail.com | (267) 296-0805
LinkedIn | GitHub | Google Scholar | Website

SUMMARY

AI security researcher and PhD candidate in ECE specializing in deepfake detection, synthetic media forensics, and adversarial ML. 5+ years of research experience producing 5 first-author publications at top-tier venues (NeurIPS, CVPR) with 60+ citations. Developed real-time detection systems achieving 97.7%+ AUC in collaboration with NVIDIA Research, securing a $150K research gift. Built and released an 8M+ frame benchmark dataset. Contributed to $3M+ in federally funded research (NSF, DARPA, NIST). Active reviewer for NeurIPS, CVPR, ICCV, and IEEE TIFS.

EDUCATION

Drexel University*Philadelphia, PA*
Ph.D., Electrical and Computer Engineering **|** *GPA: 3.85 Jan 2021 – Sept 2026 (Expected)*
Advisor: Prof. Matthew C. Stamm | Multimedia & Information Security Lab
Research Focus: Deepfake detection, synthetic media forensics, AI security, adversarial machine learning
M.S., Electrical and Computer Engineering **|** *GPA: 3.53Jan 2022 – Jan 2024*

Imam Khomeini International University*Qazvin, Iran*
B.S., Electrical Engineering **|** *Top 1% of graduating class2016 – 2020*

RESEARCH EXPERIENCE

Graduate Research Assistant*Jan 2021 – Present*
Drexel University, Multimedia & Information Security Lab | Advisor: Prof. Matthew C. StammPhiladelphia, PA
Conducting AI security research with industry collaboration with NVIDIA Research. Contributing to federally funded programs (NSF $3M+, DARPA, NIST).

Real-Time Puppeteering Defense in AI Videoconferencing | Collaboration with NVIDIA Research

- Designed the first enrollment-free defense against identity puppeteering attacks in AI-based videoconferencing systems, addressing a critical security gap in neural talking-head pipelines.
- Developed a novel pose-conditioned contrastive loss (PC-LMCL) that exploits biometric leakage in pose and expression vectors to detect identity mismatches without requiring enrollment data.
- Achieved 97.7% AUC at 75 FPS with 46% error reduction over prior state-of-the-art. Demonstrated cross-domain generalization to unseen systems (AUC 0.925).
- Work secured a $150K NVIDIA research gift and 1 year of dedicated compute resources. Published at **NeurIPS 2025** (main conference).

AI-Generated Video Detection Pipeline

- Engineered an end-to-end detection pipeline using CNNs and temporal forensic analysis to identify AI-generated videos with 99%+ accuracy.
- Validated detection across 4 major generator families: Sora, Runway AI, Stable Video Diffusion, and NeRF-based systems.
- Constructed and released a large-scale benchmark dataset of 8M+ frames spanning Transformers, Diffusion Models, and NeRFs. Hosted on Hugging Face with 400+ community downloads.
- Published at **CVPR 2024**. Dataset adopted by external research groups for benchmarking.

Transferable Detection via Virtual Generators

- Introduced a novel technique to improve synthetic video detector generalization by training with virtual generators—parametric models that synthesize forensic microstructures not associated with any real generator.
- Modeled microstructures using 2D autoregressive processes with parameters informed by architectural patterns in modern video generators.
- Demonstrated significant improvement over traditional data augmentation in generalizing to new, unseen generator architectures. Under review at **IHMMSEC 2026**.

Low-Bandwidth Talking Head Puppeteering Defense

- Proposed the first defense against puppeteering attacks in low-bandwidth talking head videoconferencing by exploiting biometric leakage in pose/expression vectors.
- Achieved 98.03% detection accuracy across 4 distinct talking head systems, flagging identity mismatches in real time. Published at **CVPR 2023 Workshop**.

GAN-Generated Face Detection via Semantic Inconsistencies

- Built a GAN-generated face detector leveraging facial semantic inconsistencies across eyes, mouth, and hair regions using feature-specific CNNs fused via SVM.
- Achieved 91.36% accuracy with demonstrated robustness to JPEG compression and resizing attacks. Published at **Electronic Imaging 2023**.

PUBLICATIONS

All first-author | 60+ citations | Cited by leading researchers including Hani Farid

- **D.S. Vahdati, T.D. Nguyen, E. Prashnani, K. Nagano, O. Gallo, M. Stamm.** “Unmasking Puppeteers: Leveraging Biometric Leakage to Expose Impersonation in AI-Based Videoconferencing.” **NeurIPS 2025.**
- **D.S. Vahdati, T.D. Nguyen, A. Azizpour, M.C. Stamm.** “Beyond Deepfake Images: Detecting AI-Generated Videos.” **CVPR 2024.**
- **D.S. Vahdati, T.D. Nguyen, M.C. Stamm.** “Defending Low-Bandwidth Talking Head Videoconferencing Systems from Real-Time Puppeteering Attacks.” **CVPR 2023 Workshop.**
- **D.S. Vahdati, M.C. Stamm.** “Detecting GAN-Generated Synthetic Images Using Semantic Inconsistencies.” **Electronic Imaging, 2023.**
- **D.S. Vahdati et al.** “Seeing the Unseen: Enhancing Synthetic Video Detector Transferability via Virtual Generators.” **Under review, IHMMSEC 2026.**

TECHNICAL SKILLS

Languages: Python, C++, C, MATLAB, Bash
ML/DL Frameworks: PyTorch, PyTorch Lightning, TensorFlow, CUDA
Deep Learning: CNNs, Transformers, GANs, Diffusion Models, NeRFs, Self-Supervised Learning, Contrastive Learning
Computer Vision: OpenCV, MediaPipe, Video/Image Forensics, Object Detection
Tools & Platforms: Git, Docker, AWS, Weights & Biases, Linux, Hugging Face

PROFESSIONAL SERVICE & LEADERSHIP

- **Peer Reviewer:** 50+ reviews for NeurIPS, CVPR, ICCV, IEEE TIFS, IHMMSEC, ICIP
- **Mentorship:** Onboarded and mentored 2 junior PhD researchers and 1 undergraduate student in the lab
- **Open-Source:** Released large-scale synthetic media dataset on Hugging Face (400+ downloads, used by external research groups)
- **Professional Membership:** IEEE