Asif Rahman

PhD Student, Department of Computer Science
1700 Hawthorne St., Apt 333, El Paso, TX 79902

arahman3@miners.utep.edu · 915-505-0596 · ucchol.github.io
Google Scholar · ORCID · DBLP · GitHub · LinkedIn

Research Interests

I work at the intersection of game theory, cognitive science, and machine learning, around a single organizing question. How should a defender plan when the adversary is a bounded human decision maker rather than a perfect optimizer, and increasingly when the adversary is an autonomous AI agent?

Classical security games assume the attacker best responds to the defender's committed strategy. Real attackers satisfice under time pressure, ignore attributes a normative model says should matter, and break down in predictable ways when a decision carries too much information. I treat those departures as the object of study. On the theory side I develop equilibrium concepts that stay useful when the defender's model of the attacker is misspecified. On the empirical side I run human subjects experiments in cyber tasks and fit process models to the resulting choices, so the behavioral assumptions inside the theory are estimated rather than assumed. On the systems side I build detectors evaluated against threats they were never trained on. Looking ahead, I am most interested in multilayer deception dynamics, adversaries that are AI agents, and defense calibrated to a distribution over attacker types rather than to a point estimate.

Keywords. Security games, Stackelberg and quantal Stackelberg equilibrium, cyber deception, bounded rationality, discrete choice and accumulator models, human subjects experiments in cybersecurity, out-of-distribution malware detection, large language models for threat detection, cyber-physical systems.

Education

The University of Texas at El Paso expected Spring 2028
PhD in Computer Science. Advisor, Christopher D. Kiekintveld.

The University of Texas at El Paso 2023
MS in Computer Science

Khulna University of Engineering and Technology, Bangladesh 2016
BSc in Electronics and Communication Engineering

Research Funding

My doctoral research is supported by the following awards to the University of Texas at El Paso. I contribute as a graduate research assistant on both.

Intelligence Advanced Research Projects Activity (IARPA) and the Office of the Director of National Intelligence (ODNI), Reimagining Security with Cyberpsychology-Informed Network Defenses (ReSCIND) program 2024–present
Award N66001-24-C-4504

DEVCOM Army Research Laboratory, Cooperative Agreement 2023–present
Award W911NF-23-2-0012

The views and conclusions in my published work are my own and do not represent the official policies, either expressed or implied, of IARPA, ODNI, the Army Research Laboratory, or the U.S. Government.

Publications

Author names are listed in publication order. My Google Scholar profile is here.

Peer-Reviewed Conference Papers

[C4] A. Rahman, A. A. N. Ryen, A. Hemida. *Beyond Best Response: Quantal Stackelberg Deception as Insurance Against Attacker Misspecification*. Conference on Decision and Game Theory for Security (GameSec), 2026. **Accepted, oral presentation.**

[C3] M. A. Sayed, A. Rahman, A. Hemida, C. D. Kiekintveld, C. A. Kamhoua. *Coordinated Multi-Domain Deception: A Stackelberg Game Approach*. International Conference on Computing, Networking and Communications (ICNC), 2026, pp. 85–89. arXiv:2601.02596.

[C2] T. Ige, C. D. Kiekintveld, A. Piplai, A. Rahman, O. Kolade, S. Kunapuli. *MAD-OOD: A Deep Learning Cluster-Driven Framework for an Out-of-Distribution Malware Detection and Classification*. IEEE International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom), 2025, Guiyang, China. arXiv:2512.17594.

[C1] A. Rahman, C. Natividad, M. A. Sayed, P. Aggarwal, C. D. Kiekintveld. *Using Choice Overload to Degrade Cyber Attacks*. Conference on Decision and Game Theory for Security (GameSec), 2025. Lecture Notes in Computer Science, vol. 16224, Springer, pp. 23–42. doi:10.1007/978-3-032-08067-7_2. Code, github.com/ucchol.

Peer-Reviewed Workshop Papers

[W1] M. A. Sayed, A. Rahman, C. D. Kiekintveld, S. García. *Fine-tuning Large Language Models for DGA and DNS Exfiltration Detection*. Annual Computer Security Applications Conference Workshops (ACSAC Workshops), 2024. doi:10.1109/ACSACW65225.2024.00034. arXiv:2410.21723.

Manuscripts in Preparation

[M1] A. Rahman et al. *Process models of adversarial choice under information load*. Cognitive process modeling of attacker target selection. In preparation for CHI 2027.

[M2] A. Rahman et al. *Multilayer deception dynamics against agentic AI attackers*. Deception policy across coupled cyber and physical layers with a learning adversary. In preparation for AAMAS 2027.

Research Projects

1. Adaptive AI-Based Cybersecurity Detection 2024–present

Most malware and command-and-control detectors are trained and scored under a closed world assumption, so their reported accuracy says nothing about the first day a genuinely new threat arrives. This project builds detectors that are evaluated only on what they have never seen. One line fine-tunes large language models to flag domain generation algorithms and DNS exfiltration, tested across 59 real-world DGA families with families held out of training rather than shuffled into it. A second line, MAD-OOD, models each malware family as a class-conditional spherical decision boundary derived from Gaussian discriminant analysis, then uses z-score distance across class centroids to separate in-distribution from out-of-distribution samples with no OOD data at training time. A second stage network fuses cluster assignments, refined embeddings, and supervised classifier outputs, reaching an AUC near 0.91 on unseen families.

Papers, [W1] arXiv:2410.21723 and [C2] arXiv:2512.17594.
Code, github.com/ucchol/nxdomain-sentinel.

2. Human-Centered Cybersecurity Using Cognitive Science and Decision-Complexity Frameworks 2023–present

If attackers are people, then the limits of human choice are a defensive resource. This project designs and runs behavioral experiments in which participants perform an adversarial file selection task alongside a matched consumer purchasing task, under a within-subjects factorial manipulation of how many options are shown and how much attribute detail each option carries. I fit conditional logit, mixed logit, multi-attribute linear ballistic accumulator, and ordinal comparison models to the resulting choices. The central result is that attribute detail rather than option count is what breaks the link between a person's stated preferences and their observed choice, and that the collapse is invisible to the measures practitioners usually trust, since outcome quality and self-reported difficulty both fail to track it. The project also produced a general complexity and time framework in Python that formalizes comparison work, a speed-accuracy knee point, an overload boundary, and the rational switch from compensatory weighing to a

single-cue heuristic.

Papers, [C1] doi:10.1007/978-3-032-08067-7_2 and [M1] in preparation.
Code, github.com/ucchol/Choice-Overload-in-Cyber-Defense.

3. AI-Enabled Security Through Game-Theoretic Optimization 2023–present

Strong Stackelberg equilibrium tells a defender how to commit against an attacker who best responds exactly. Get the attacker model slightly wrong and that guarantee can evaporate. This project replaces the perfectly rational attacker with a quantal Stackelberg equilibrium built on logit choice, and asks what the defender gives up and what it gains. I built the full research stack from scratch, including a gradient-based QSE solver with an analytical Jacobian and multi-start restarts, reward matrix builders, eight experiment scripts, and a 16-node cyber-physical testbed drawn from real CVE records with CVSS-derived exploit probabilities. The main theoretical result is an insurance property, proved for any finite Stackelberg game with logit response. Quantal deception concedes very little value when the attacker model is right and protects substantial value when it is wrong, with the gap traced to a tie tax that separates booked from realized defender payoff on a tie boundary. A companion effort extends the same machinery to deception coordinated across coupled cyber and physical layers.

Papers, [C4] GameSec 2026 and [C3] arXiv:2601.02596.
Code, github.com/ucchol/qse-cyber-deception.

Projects

Public research code, released under an MIT license with reproducible notebooks and committed result files.

network-defense-rl 2026

A Gymnasium environment and agent suite for sequential network defense, where a defender allocates scarce hardening and incident response across a network whose size changes every episode while a boundedly rational attacker probes it. I trained PPO and DQN defenders and benchmarked them against a quantal Stackelberg equilibrium solver rather than against weak baselines, and the trained defender tracks that solver on network sizes it never saw during training. I measured where the policy places its actions to establish size generalization directly, swept attacker rationality across two orders of magnitude to test how much the defense depends on knowing the adversary, and isolated the contribution of the equilibrium solve with a paired test against simpler value ordered allocation.

Code, github.com/ucchol/network-defense-rl. Gymnasium, Stable-Baselines3, PyTorch, SciPy.

nxdomain-sentinel 2026

A domain threat classification service scored only on malware families it has never seen. In place of the random shuffle behind most reported DGA accuracy, the evaluation holds every family out in exactly one of five folds and trains it in the other four, so each family is compared against itself and family difficulty cancels, which quantifies what generalization to a new family actually costs. I added a novelty gate that fits one Gaussian per known family in the classifier penultimate space with a pooled shrunk covariance and scores Mahalanobis distance to the nearest centroid, requiring no out-of-distribution data to fit, and it fires on the families the binary classifier has no representation for. The service is containerized and serves in about one millisecond per domain at batch size 100.

Code, github.com/ucchol/nxdomain-sentinel. PyTorch, FastAPI, Docker, Cloud Run.

Honors and Awards

Partial travel grant, GameSec 2025, Athens, Greece 2025

Research Experience

Graduate Research Assistant, The University of Texas at El Paso 2021–present
Department of Computer Science. Advisor, Christopher D. Kiekintveld.

- Built the research infrastructure for a quantal Stackelberg approach to cyber deception, including a gradient-based solver with an analytical Jacobian and multi-start restarts, reward matrix builders, eight experiment scripts, and a 16-node cyber-physical testbed drawn from real CVE and CVSS records. Proved that the resulting defense behaves as insurance, giving up little value when the attacker model is correct and protecting substantial value when it is misspecified [C4].
- Established the theoretical generality of that result, restating and proving the main theorem for any finite Stackelberg game with logit response, adding an existence result that needs no tie-breaking convention, and formalizing the tie tax as the gap between booked and realized defender value on a tie boundary. Validated every reported number against regenerated ground truth matrices.
- Designed and ran a behavioral experiment on adversarial file selection paired with a consumer purchasing control task, then modeled the data with conditional logit, mixed logit, and multi-attribute accumulator models. Attribute detail rather than option count is what breaks the link between stated preferences and observed choice, and the effect is invisible to outcome quality and to participants' own self-reports.
- Built a reusable complexity and time framework in Python that formalizes comparison work, a speed-accuracy knee point, an overload boundary, and the rational switch from compensatory strategies to single-cue heuristics, then applied it across two decision domains.
- Developed coordinated deception across coupled cyber and physical layers, using CVSS-derived exploit probabilities to place deceptive resources under a Stackelberg formulation [C3].
- Fine-tuned large language models to detect domain generation algorithms and DNS exfiltration, evaluated across 59 real-world DGA families with emphasis on generalization to families withheld from training [W1].
- Contributed to MAD-OOD, a two-stage cluster-driven framework for out-of-distribution malware detection that models family embeddings with class-conditional spherical decision boundaries and requires no out-of-distribution data during training [C2].
- Released the research code behind this work publicly, including a quantal Stackelberg solver, a variable-size network defense environment whose agents are benchmarked against that solver, and a containerized domain threat detection service evaluated by family holdout.

Presentations

Using Choice Overload to Degrade Cyber Attacks. GameSec 2025, Athens, Greece 2025

Teaching Experience

Graduate Teaching Assistant, The University of Texas at El Paso 2021–present
Department of Computer Science. Courses supported include Operating Systems, Computer Networks, and Machine Learning. Led laboratory sections, held weekly office hours, graded assignments and exams, and mentored student course projects through design, implementation, and final presentation.

Professional Service

Program Committees

Program Committee, AI for Social Impact (AISI) track, AAAI 2027 2026
Program Committee, AI for Social Impact (AISI) track, AAAI 2026 2025

Reviewing

Reviewer, International Conference on Computing, Networking and Communications (ICNC) 2026
Reviewer, International Conference on Electrical, Computer and Energy Research (ICECER) 2026
Reviewer, International Conference on Autonomous Agents and Multiagent Systems (AAMAS) 2025

Technical Skills

Languages	Python, C++, Java, JavaScript, Bash
Modeling and ML	PyTorch, TensorFlow, scikit-learn, NumPy, SciPy, pandas. LLM fine-tuning with LoRA, discrete choice and accumulator models (conditional logit, mixed logit, MLBA), constrained and gradient-based optimization, clustering, Gaussian discriminant analysis
Reinforcement Learning	Gymnasium environment design, Stable-Baselines3, PPO and DQN, domain randomization, shared-seed evaluation protocols
Game Theory	Stackelberg and quantal Stackelberg solvers, security game formulation, reward matrix construction, equilibrium computation and validation
Behavioral Methods	Human subjects experiment design, within-subjects factorial designs, Prolific recruitment, response time analysis, model comparison by log-likelihood, BIC, and held-out transfer
Security	Game-theoretic deception modeling, CVE, CVSS, and NVD vulnerability data, malware and DGA detection, out-of-distribution evaluation, packet analysis
Systems and Tools	Linux, Git, Docker, Kubernetes, AWS, Google Cloud Run, FastAPI, PostgreSQL, MySQL, LaTeX

Professional Experience

bKash Limited, Bangladesh 2018–2021

Engineer, Service Operations. Mobile financial services platform serving in-app banking at national scale.

- Ran a hybrid production environment spanning AWS and on-premises infrastructure, covering server provisioning, service configuration, and routine capacity and health checks across both sides of the stack.
- Handled Linux system administration in both environments, including operating system and package patching, user and access management, scheduled backups, cron jobs, and log rotation on production hosts.
- Owned incident triage for production issues, working from the first alert through diagnosis and mitigation to engineering handoff, and traced recurring failure modes back to their source rather than closing them one at a time.
- Integrated and supported payment APIs between the banking platform and partner systems, validating request and response flows and debugging transaction failures reported by internal and merchant-facing teams.
- Automated log parsing and operational reporting in Python, replacing recurring manual reporting with scheduled jobs and shortening the turnaround on routine data requests.
- Supported change and release management for production deployments, coordinating maintenance windows with engineering and business stakeholders and verifying service health after each rollout.

REVE Systems Limited, Bangladesh 2016–2018

System Administrator. VoIP and telecommunications software provider.

- Administered the production Linux servers underpinning VoIP services, covering installation, configuration, patching, user and permission management, and scheduled backup routines.
- Maintained production MySQL servers, including user and schema administration, backup and restore procedures, replication health checks, and query performance troubleshooting.
- Diagnosed VoIP service disruptions by capturing and analyzing SIP signaling and UDP media traffic, isolating faults between the platform, the internal network, and upstream carriers.

- Monitored service availability and system resource usage, responded to alerts outside business hours, and escalated to engineering with the packet captures and log evidence needed to reproduce a fault.
- Applied configuration changes and service updates across production nodes during scheduled maintenance windows, with rollback steps prepared in advance.
- Documented recurring faults and their resolution paths so that repeat incidents could be handled by the operations team without escalation.