MITCHELL OLDHAM

mwo@kralizec.com | mitchelloldham.com | mitchell-oldham | TheWhiteTower16

EXPERIENCE

AI and Security Consultant *May 2025 - present*
Kralizec
- Architected secure AI agent solutions for 20+ companies, increasing operational efficiency while reducing vulnerabilities.
- Designed and delivered over 100 custom benchmarks that help clients optimize their AI systems for domain-specific tasks.
- Advised 5+ organizations on deploying and maintaining self-hosted AI models for data providence and governance.
- Overseen 10+ enterprise AI adoption programs, guiding them on where and how to deploy AI effectively.
- Aided partners on integrating AI security tools, helping decrease incident response time and increase preparedness.
- Built custom MCPs, RAG databases, fine-tuning systems, and RL pipelines for enterprise AI deployment customization.
- Identified and reported safety and security relevant risks in frontier AI systems to their respective providers.

Software Engineering Intern *May 2024 - Feb 2025*
Johnson Controls
- Developed API endpoints, standardized error responses, and automated test workflows leading to a 6% uptime increase.
- Co-designed and implemented an object creation system that bridged frontend services to backend databases.
- Optimized context engineering for certain OpenBlue genAI features, reducing both cost and inference time by 20%.
- Assisted in the deployment of internal AI developer tools, which achieved over 80% adoption among targeted teams.

System Security Researcher *May 2023 - Dec 2023*
WashU McKelvey School of Engineering
- Discovered 11 vulnerabilities in Microsoft's GPT-4 powered Nuance DAX using adversarial prompting and jailbreaking.
- Refactored ROS sensor drivers, cutting robot startup time by 41% and memory consumption by 25%.
- Uncovered emerging security risks by analyzing adoption trends of AI-enhanced tooling across several key industries.
- Secured IoT and network infrastructure for organizations around St. Louis by testing systems and fixing vulnerabilities.

Oncology Machine Learning Engineer *May 2022 - Dec 2022*
University of Missouri–St. Louis
- Improved KNN, SVM, LR, and RF models, increasing average prediction accuracy from 94.1% to 98.3%.
- Identified new quinazoline-based MET inhibitors with 9% higher binding efficacy when compared to existing drugs.
- Validated computational findings by coordinating the testing of drug candidates in wet-lab experiments.
- Strengthened security for our web tool EDock-ML by implementing a file validation protocol, preventing upload attacks.

PROJECTS

- Conducted pre-launch red-teaming on several OpenAI systems, including GPT-5 and ChatGPT Agent
- Built a multi-agent penetration testing framework that orchestrates the collection, scanning, and exploitation of a target.
- Worked as an external tester for Anthropic, focused on Constitutional Classifiers and agent-level security measures.
- Reported vulnerabilities in popular AI applications including Claude Code, Microsoft Copilot, and Salesforce Einstein.
- Ranked in the top 10 in the Gray Swan Agent Red-Teaming Arena by bypassing the safeguards of 21 AI-powered systems.
- Trained a local cybersecurity AI that helps with threat hunting, code analysis, reverse engineering, and malware analysis.

EDUCATION

Washington University in St. Louis *GPA: 3.9/4.0*
Bachelor of Science in Mathematics and Computer Science
Courses: Linear Algebra, Machine Learning, Quantum Information Processing, Cryptography, Cloud Computing, Computer Networks, Advanced Algorithms, Probability, Numerical Methods, Data Science, Statistics

SKILLS

Languages:	Python, Java, Rust, C#, SQL, JavaScript, R, MATLAB, BASH, Qiskit
Tools:	Linux, VMware, Azure, AWS, Splunk, Docker, Wireshark, Git, Kubernetes, PostgreSQL, PyTorch
Experience:	Vulnerability Assessment, IAM, Compliance, Threat Modeling, Agile, CI/CD, RAG, MCP
Standards:	ISO 27001, MITRE ATLAS, NIST RMF, OWASP Top 10, HIPAA