Alexander Nemecek

alex.nemecek@att.net | (216) 501-3231 | LinkedIn | Personal Site

EDUCATION

Case Western Reserve University – Cleveland, OH Aug 2023 – Expected Aug 2027
Ph.D. Candidate, Computer Science
Focus: Watermarking in Large Language Models (LLMs), AI Governance

Ohio University – Athens, OH Aug 2019 – May 2023
B.S., Computer Science – B.S., Psychology

TECHNICAL SKILLS

Languages: Python, C/C++, JavaScript, TypeScript, SQL, Bash
ML/LLMs: PyTorch, TensorFlow, Hugging Face Transformers, PEFT/LoRA, supervised fine-tuning, RAG, LLM evaluation & benchmarking
Agentic AI: LangChain, LangGraph, tool/function calling
Tools & Platforms: AWS, Git, Linux, Slurm, Snowflake, Pandas, NumPy, Tableau, React Native

EXPERIENCE

Advanced Micro Devices (AMD) – San Jose, CA Sep 2026 – Present
Artificial Intelligence Framework Intern

- Fall internship on AMD's model optimization team: LLM quantization with Quark, AMD's open-source quantization framework, and inference optimization techniques such as AWQ, GPTQ, and speculative decoding.

Case Western Reserve University – Cleveland, OH Aug 2023 – Present
Graduate Researcher

- Developed semantic-aware watermarking scheme for LLMs using topic-aligned token partitioning in PyTorch, maintaining text quality at industry-leading levels and **increasing adversarial robustness by 30%**.
- Directed cross-disciplinary AI Governance analysis, designing a three-layer policy-technical framework adopted to align regulatory expectations with watermarking capabilities.
- Engineered and deployed end-to-end AI stress detection platform (AWS Lambda, Fitbit API, React) with personalized models; **led an Agile team of 10 junior engineers**.

Procter & Gamble – Cincinnati, OH May 2026 – Aug 2026
Data Science Intern

- Architected and deployed agentic hazard-assessment pipeline (LangChain, LangGraph) orchestrating tools, data retrieval, and reasoning per P&G decision methodologies, reducing assessment **from 6 weeks to 1-2 days**.
- Integrated vetted external toxicology data sources via custom FastAPI endpoints, replacing manual evidence gathering with automated live retrieval.
- Benchmarked and implemented structural-alert models classifying chemical hazards from SMILES data, **achieving 80% accuracy standalone and 90% with ensembles**.

Sandia National Laboratories – Albuquerque, NM May 2024 – Mar 2026
Research and Development Intern

- Built real-time threat detection pipeline leveraging fine-tuned LLMs with LoRA in PyTorch, **improving model accuracy by 20%** and enabling production deployment.
- Implemented adaptive LoRA fine-tuning infrastructure for dynamic model swapping, speeding up knowledge graph extraction workflows.
- Designed a prompt engineering framework with semantic guardrails; **adopted by 3 security teams** for strategic threat prioritization.

Cisco Systems – Richfield, OH May 2023 – Aug 2023
Software Engineer Intern

- Rebuilt full-stack BOM analysis platform (Python, Snowflake, React), automating manual processes into a user-friendly application and transforming hardware engineering workflows from error-prone to efficient.
- Developed Python data pipeline standardizing vendor BOM formatting into Snowflake warehouse with integrated SQL pipelines, **reducing analysis time by 60%**.
- Created interactive pricing trend visualizations in React that **accelerated procurement decision cycles from 3 months to 1 month** across cross-functional teams.

Sherwin-Williams – Cleveland, OH May 2022 – Aug 2022
Software Engineer Intern

- Constructed SQL and Snowflake pipelines to centralize nationwide automotive paint data, improving accessibility and operational efficiency.

- Designed Tableau analytics platform **visualizing 20+ KPIs**, resolving data inconsistencies to improve inventory management and sales forecasting decisions.

Ohio University – Athens, OH Sep 2021 – Mar 2023
Student Researcher

- Implemented secure cross-platform mobile app for foster care memory preservation with AWS backend storage and privacy-compliant architecture, integrating CI/CD pipeline with automated testing.
- Developed VR/AR control interfaces for industrial robotics using Meta Quest and HoloLens, improving remote worker safety operations.

PUBLICATIONS & PREPRINTS

- **A. Nemecek**, E. Yilmaz, E. Ayday, "Cluster-Aware Attacks on Graph Watermarks," *ESORICS*, Sep 2026.
- O. Zafar*, **A. Nemecek***, E. Ayday, "The End of Trust: How Agentic AI Breaks Security Assumptions," *NSPW*, Sep 2026. **Equal Contribution*
- H. He, S. Xu, **A. Nemecek**, J. Li, E. Ayday, G. Cheng, "Optimal Watermark Generation under Type I and Type II Errors," *Biometrika*, Aug 2026.
- **A. Nemecek**, O. Zafar, D. Ganguly, V. Singh, V. Chaudhary, E. Ayday, "Auditing Cross-Lingual Fairness in Language Model Watermarking," *arXiv*, Aug 2026. *Under review at NeurIPS 2026*.
- Y. Xu, O. Zafar, **A. Nemecek**, E. Ayday, "Beyond Detection: Evaluating Defensive LLMs Against AI-Generated Social Engineering in Live Turn-by-Turn Interaction," *arXiv*, Aug 2026. *Under review at AAAI 2027*.
- **A. Nemecek**, Y. Jiang, E. Ayday, "Topic-Based Watermarks for Large Language Models," *ACL*, Jul 2026.
- **A. Nemecek**, W. Li, X. Jiang, J. Vaidya, E. Ayday, "Quantifying Memorization and Privacy Risks in Genomic Language Models," *ACM-BCB*, Jul 2026.
- **A. Nemecek**, H. He, G. Cheng, E. Ayday, "Authenticated Contradictions from Desynchronized Provenance and Watermarking," *CVPR APAI Workshop*, Jun 2026. ***Best Paper Award***
- **A. Nemecek**, O. Zafar, Y. Xu, W. Li, E. Ayday, "Who Gets Flagged? The Pluralistic Evaluation Gap in AI Content Watermarking," *CVPR MAPS Workshop*, Jun 2026.
- O. Zafar, **A. Nemecek**, Y. Zhang, W. Li, D. Ganguly, V. Singh, V. Chaudhary, E. Ayday, "Privacy Policy Enforcement Guardrails for Data-Sensitive Retrieval-Augmented Generation," *arXiv*, May 2026. *Under review at TPS 2026*.
- **A. Nemecek**, Y. Jiang, E. Ayday, "Watermarking Without Standards Is Not AI Governance," *arXiv*, Mar 2026. *Under review at NeurIPS 2026*.
- **A. Nemecek**, Y. Jiang, E. Ayday, "The Feasibility of Topic-Based Watermarking on Academic Peer Reviews," *IJCNLP-AACL*, Dec 2025.
- M. Namazi, **A. Nemecek**, E. Ayday, "ZKPROV: A Zero-Knowledge Approach to Dataset Provenance for Large Language Models," *arXiv*, Dec 2025. *Under review at PoPETS 2027*.
- **A. Nemecek**, Z. Yun, Z. Rahmani, Y. Harel, V. Chaudhary, M. Sharif, E. Ayday, "Exploring Membership Inference Vulnerabilities in Clinical Large Language Models," *IEEE TPS-ISA HMD-SPiRiT Workshop*, Nov 2025.

PRESENTATIONS

- **A. Nemecek**, R. Abbott, C. Garasi, "Can Smaller Expert Modules Enhance RAG Performance?," *Chesapeake Large-Scale Analytics Conference (CLSAC)*, Nov 2024.
- **A. Nemecek**, Y. Jiang, E. Ayday, "Topic-Based Watermarks for Large Language Models," *Privacy Enhancing Technologies Symposium (PETS)*, Jul 2024.
- **A. Nemecek**, C. Mourning, "Detection of Network Attacks using Machine Learning Models," *The Ohio Journal of Science*, Apr 2023.