Adrian Lara

contact@adrianlara.dev | linkedin.com/in/adrian-lara-cs | github.com/adlune | adrianlara.dev

Education

National Autonomous University of Mexico (UNAM) Morelia, Michoacan

Bachelor of Science in Information Technologies for Science *2024 – Present*

- Relevant Coursework: Machine Learning, Neural Networks, Data Mining, Inferential Statistics, Multivariate Statistics, Modeling and Simulation, Distributed Databases, Networks and Telecommunications, Linear Algebra

Projects

Force-Directed Network Anomaly Detection | *Python, NumPy, pandas, Matplotlib, pytest* 2026

Modeling and Simulation course project, UNAM github.com/adlune/force-directed-NBAD

- Built an N-body simulation that turns each minute of network traffic (CICIDS2017) into a graph of hosts and relaxes it to physical equilibrium, testing its final energy as a novelty-detection signal against benign traffic.
- Fixed divergence in the standard Fruchterman-Reingold update (errors grew 5× per step) with per-node Jacobi preconditioning, proven stable by Gershgorin's theorem. Halved per-iteration cost and parallelized across windows (2.8× on 4 cores).
- Showed that the benchmark's NAT-routed DDoS and port scans leave the host graph unchanged (identical in all 21 DDoS windows), then designed an attack-injection study instead: 774 simulations, 3 attack shapes, 6 sizes.
- Found that the energy flags attacks only above about 80 new connections and that four plain graph counts predict it ($R^2 = 0.9994$), so simple counts detect earlier. Backed by 30 unit tests and reproducible figures.

Groundwork RAG | *Python, FastAPI, ChromaDB, sentence-transformers, Ollama, pytest* 2026

Information Systems and Web Services course project, UNAM github.com/adlune/groundwork-rag

- Built a local retrieval-augmented generation app without RAG frameworks: PDF cleanup, chunking, bge-small embeddings in ChromaDB, a local LLM through Ollama, and a FastAPI web app that cites sources and shows each retrieved passage with its scores.
- Added BM25 keyword search merged through Reciprocal Rank Fusion, raising hit@5 from 0.71 to 0.79 on a 45-question evaluation set and from 0.50 to 1.00 on short single-term questions.
- Calibrated a two-threshold refusal gate with no wrong refusals or off-topic answers on the evaluation set (0.2% and 0.7% on held-out halves across 300 splits). Evaluation runs are fully reproducible, backed by 64 unit tests.

Crosswire IDS | *Python, Polars, pandas, Matplotlib, Jupyter* 2026

Personal project, in progress github.com/adlune/crosswire-ids

- Building an IoT intrusion detection system trained on CICIoT2023 and tested on traffic from a virtual IoT lab it has never seen, to measure how far detection drops, for which attacks, and because of which features.
- Found in EDA that the dataset's window size leaks the label (a rule with no learning reaches 97.4% accuracy), that 55% of rows are exact duplicates, and that features such as TTL describe the lab rather than the attacks.

Cybersecurity Lab & Notes | *Kali Linux, Hack The Box* 2025 – Present

Self-study github.com/adlune/notes

- Practice offensive and defensive security on Hack The Box machines from a self-managed Kali Linux VM.
- Document write-ups and notes on machines and cybersecurity concepts, linking each concept to practical cases.

Skills

Programming: Python, SQL, R, Bash, HTML/CSS
Data & ML: NumPy, pandas, Polars, Matplotlib, scikit-learn, PyTorch, Jupyter
LLMs & Retrieval: RAG, sentence-transformers, ChromaDB, BM25, Ollama
Development: FastAPI, pytest, Git, uv, Docker, Linux (Kali/Ubuntu/Arch)
Security: Offensive and defensive security, network traffic analysis, intrusion detection, Kali Linux
Languages: Spanish (native), English (TOEFL ITP: 637/677, CEFR C1 – Proficient User)