YAIMA VALDIVIA

Principal Software Engineer

650-484-9840 yaimamvaldivia@gmail.com yaimavaldivia.com linkedin.com/in/yaimavaldivia

SUMMARY

Principal Software Engineer who builds AI agent systems and the security layer that governs them. Creator of Agent Trust (open-source runtime authorization for agents across MCP, A2A, and API) and HIVE (prompt-injection detection), and builder of a multi-agent platform for autonomous task coordination. 10+ years in backend and security engineering. MS in Computer Science, AI specialization, Georgia Tech

EXPERIENCE

05/2026 - Present

Principal Software Engineer

NetScout

- Engineer DDoS detection and mitigation systems for large-scale network security

01/2025 - 05/2026

Principal Software Engineer

DigiCert

- Built an LLM-powered analytics tool that turns statistical analysis into user-friendly reports with natural-language explanations (internal hackathon project)
- Built ML-based DDoS detection using Random Forest classification on the CIC-DDoS2019 dataset, reaching 99% accuracy (internal hackathon project)
- Built real-time proxy traffic monitoring on ClickHouse
- Designed Splunk HEC telemetry integration with 93% test coverage
- Built and operated security services in Go with PostgreSQL, MongoDB, and GraphQL

05/2021 - Present

Instructional Associate

Georgia Institute of Technology

- Instructional Associate for Game Artificial Intelligence and Video Game Design, supporting student learning in applied AI

12/2021 - 01/2025

Principal Software Engineer

Vercara

- Built cybersecurity services in Go with PostgreSQL, MongoDB, and GraphQL
- Contributed to AI research and applied-ML initiatives across security products

11/2020 - 11/2021

Principal Software Engineer

Neustar

- Built security services in Go with PostgreSQL, MongoDB, and GraphQL
- IoT: developed an IoT platform adhering to Open Connectivity Foundation (OCF) specifications to address interoperability, including a PKI-like solution for Eclipse IoT projects (SmartHome, Kura, Hawkbit) and Android applications. Shared industry knowledge through public speaking

05/2016 - 11/2020

Senior Software Engineer

Neustar

- Built security services in Go with PostgreSQL, MongoDB, and GraphQL

09/2014 - 04/2016

Software Development Engineer

Samsung Electronics America

- Integrated connected devices into the SmartThings ecosystem and supported the developer community
- Developed home automation solutions, liaising with device manufacturers for integration
- Supported key industry events including CES. Founded Women in IoT

05/2014 - 08/2014

Analytic Consultant - Lead Consultant, Americas

FICO

- Used FICO Model Builder for data processing, validation, and analysis
- Developed credit bureau characteristics and audit materials in Java and Groovy

11/2011 - 12/2014

Lead Engineer and Founder

Vee Collective, LLC

- Led development of applications across iOS, Android, and Web, collaborating with design to ship products

EXPERIENCE

03/2012 - 03/2014

Development Analyst Senior

TransUnion

- Built custom credit application in C# and SQL using localized data
- Created custom reports and provided training in English and Spanish for international clients

05/2006 - 03/2012

Software Solution Specialist

CRIF Lending Solutions

- Built and maintained credit decisioning solutions using StrategyOne, Java, and C# for US and international clients
- Delivered technical documentation and training for production deployments

EDUCATION

2018 - 2022

Master of Science in Computer Science, Artificial Intelligence

Georgia Institute of Technology

2002 - 2004

Bachelor of Science, Computer Systems Engineering

Universidad Latina de Costa Rica

SKILLS

LLM & Agents

agent authorization **prompt-injection detection** **multi-agent orchestration** **MCP** **A2A**

managed-agent runtimes **LangChain**

Machine learning

PyTorch **HuggingFace Transformers** **sentence-transformers** **scikit-learn** **CodeBERT** **DeBERTa**

Random Forest **Thompson-sampling bandits** **PyACTR** **ACT-R**

Backend / infrastructure

Postgres **MongoDB** **ClickHouse** **Redis** **GraphQL**

Security / identity

OAuth 2.1 **SPIFFE/WIMSE** **JWT** **Ed25519 JWS** **FERPA-compliant hashing**

AI RESEARCH AND INDEPENDENT PROJECTS

Agent Trust

Open-source runtime authorization for AI agents. Issues opaque, session-bound tokens and evaluates every action with instant revocation, protocol-agnostic across MCP, A2A, and REST API

- Risk-routed Guardian pipeline: a rule-based prefilter, a policy engine, and LLM-based reasoning for destructive and admin actions, classifying actions across four effect categories (read, mutating, destructive, admin)
- Tamper-evident per-session audit chain (SHA-256 hash chaining) with SDK callbacks for LangChain, CrewAI, AutoGen, and Vercel AI
- 9 services in Go with Redis and Postgres for session state and audit history. SDKs in 5 languages. Licensed Apache 2.0

AI RESEARCH AND INDEPENDENT PROJECTS

Stage Pilot

Multi-tenant musician-management platform. Three role-specialized LLM agents (Scout, Outreach, Booking) provisioned per artist on a managed-agent runtime, each with an isolated memory store so context never crosses roles

- Per-tool trust-phase ladder (Apprentice, Trial, Autonomous) enforced inside tool handlers, providing a bypass-proof gate independent of model behavior
- Reactive coordination via Gmail Pub/Sub and Calendar push channels, scheduled wakes via Inngest, and a custom SSE session runner. 20 Prisma models, ~18,600 LOC

HIVE

Honey-Prompt Injection Detection

- Open-source, model-agnostic, real-time prompt-injection detector. An I/O overlay around any LLM with honey-token tripwires that detect context compromise even after Unicode homoglyph, zero-width, Base64, URL/HTML/hex, or ROT13 obfuscation, via output canonicalization before token matching
- Four-stage cascade (heuristic rules, embedding-based attack-memory similarity, honey-token leakage check, optional LLM judge) achieving F1 0.889 / precision 0.931 / recall 0.850 / AUC-ROC 0.908 on a 10,958-sample benchmark, matching ProtectAI DeBERTa-v3 with no task-specific training and no GPU
- 3x advantage over Rebuff on a 17-type obfuscation suite (TPR 0.706 vs 0.235); 13.4 ms per sample. 34-page manuscript submitted to ACM, May 2026

Episodic Memory in PyACTR

Cognitive-architecture research

- Integrated an episodic memory module into PyACTR (Python ACT-R cognitive architecture) for context-dependent linguistic tasks - ambiguous homonyms, rare irregular forms, novel nonsense words, suppletive forms, high-conflict edge cases. Context vector enrichment with temporal, performance, and recency features; layered retrieval tolerates 8% first-attempt accuracy via backup processing
- Episodic memory reduced average processing time ~44% (5.4 ms vs 9.6 ms baseline) at 100% task accuracy; ablation found precise matching outperforms partial matching for episodic recall, lifting success from 45-50% to 100% under the same configuration

Similarity Check

Academic integrity platform

- Multi-tenant SaaS detecting peer plagiarism, LLM-generated code, and cross-semester recycling in game-development courses. JWT auth, role-based access, per-institution white-label branding, FERPA-compliant HMAC name hashing
- Weighted score fusion across five signals (MOSS winnowing, JPlag-style tokenization, structural features, CodeBERT embeddings, text difflib); mathematical upper-bound pruning eliminates 94% of peer pairs and 99.9% of historical pairs, taking 193K comparisons from a projected 18 hours to about 3 minutes
- Verified at 622 students plus 1,092 historical submissions, with Unity archives in the 10GB range. ~22K LOC Python, ~9,500 LOC TypeScript, 328 tests

Lyric Forge

Mobile songwriting assistant. LLM-assisted editing parses the active draft, infers edit scope (phrase, section, or insertion), and conditions generation on the surrounding section to preserve the writer's voice, tense, and line structure

- Two-tier domain enforcement (lexical signal prefilter + LLM sentinel response) with tiered model routing for synonyms, rhymes, and full lyric generation. Custom 2,100-line music-theory library (12-key transposition, diatonic chords, circle of fifths) and 7 contract test suites (~1,100 lines)