

Stephen Sequenzia

Orlando, FL | stephen@sequenzia.com | 407.718.2010 | linkedin.com/in/sequenzia | github.com/sequenzia

ACTIVE SECRET SECURITY CLEARANCE

PROFESSIONAL SUMMARY

Senior Staff AI/ML Engineer and Architect with 20+ years of experience spanning production software, data, ML, and AI systems. At Lockheed Martin, lead AI/ML architecture and technical strategy for agentic AI systems across classified and unclassified defense programs, guide multiple engineering teams across 8–10 concurrent efforts, and drive agentic engineering enablement for 500+ software and cyber engineers. Hands-on work includes production agents, custom Model Context Protocol (MCP) services, evals, guardrails, Pydantic AI, LangChain/LangGraph, Kubernetes, and AWS GovCloud. Key results include a shared MLOps framework that cut deployment time by roughly 85% and PyTorch optimizations that reduced deep learning training time by over 94% on average. Previous experience includes founding and operating technology consulting and ML R&D businesses and directly designing and training CNNs, RNNs, Transformers, Bayesian neural networks, and deep ensembles across research and production systems.

TECHNICAL SKILLS

Agentic AI & Context Systems

Agentic Systems Architecture, Context Engineering & Memory, Multi-Agent Runtimes & Orchestration, MCP & Tool Integration, RAG & Retrieval Pipelines, Agent Evaluation, Guardrails & Observability, Agentic Development Workflows

Tools: Pydantic AI, LangChain, LangGraph, MCP, Pydantic Evals, Pydantic Logfire, Claude Code, Codex

Machine Learning & Applied Modeling

Machine Learning Engineering, Deep Learning, Time-Series Modeling, Statistical Modeling, Feature Engineering, Neural Architecture Design, Probabilistic & Bayesian Modeling, Model Training & Performance Optimization

Tools: PyTorch, TensorFlow, Keras, Scikit-learn, NumPy, Pandas, SciPy, statsmodels

MLOps & Data Platforms

MLOps & ML Platform Architecture, Model Lifecycle Engineering, Distributed Data Pipeline Architecture, ETL & Data Integration, Lakehouse & Feature Store Architecture, Pipeline Orchestration, Reproducibility, Lineage & Monitoring, GPU & Accelerated Computing

Tools: MLflow, Apache Airflow, Apache Arrow, Apache Parquet, Dask, Ray, Flyte, RAPIDS, NVIDIA CUDA

Software & Systems Architecture

Software Architecture, Backend & API Engineering, Full-Stack Web Systems, Framework & Library Design, Distributed & Event-Driven Systems, Real-Time Streaming Systems, Developer Tooling & Automation, Durable & Restart-Safe Workflows

Tools: Python, TypeScript, C++, Node.js, FastAPI, React, Next.js, Vercel AI SDK, DBOS

Databases & Enterprise Integration

Database Architecture & Data Modeling, SQL Engineering, SQL/NoSQL Administration, Enterprise Application Integration, Schema & Query Performance Engineering

Tools: Microsoft SQL Server, PostgreSQL, SQLite, SQLAlchemy, Alembic, WebSockets

Cloud, Infrastructure, Reliability & Security

Cloud & Hybrid Infrastructure Architecture, Networking & Systems Engineering, High Availability & Platform Reliability, Container Platform Operations, Security, Identity & Authorization, Compliance-Controlled Deployments

Tools: AWS Cloud, Microsoft Azure, Google Cloud, Kubernetes, Docker, Amazon S3

Engineering Leadership & Enablement

Technical Strategy & Architecture Leadership, Engineering & Team Leadership, Technology Consulting & Stakeholder Advisory, Business Ownership & Technical Entrepreneurship, Hiring, Mentoring & Team Development, Organization-Wide AI Enablement, Curriculum Development & Technical Instruction

Lockheed Martin Corporation

Apr 2026 - Present

Orlando, FL

Senior Staff AI/ML Engineer & Architect

Lead AI/ML architecture and technical strategy for agentic AI systems across classified and unclassified defense programs. Guide multiple engineering teams across 8–10 concurrent efforts and drive organization-wide AI enablement for 500+ software and cyber engineers, with hands-on contributions spanning context engineering, multi-agent runtimes, MCP services, evals, guardrails, shared MLOps infrastructure, and distributed data platforms.

- Set technical direction and architectural consistency across 8–10 concurrent AI/ML efforts on the engineering leadership team, and serve as a senior lead for MFC's AI Enablement initiative: developing and teaching agentic engineering curricula, leading office hours, mentoring 10+ engineers, and growing teams through technical interviews and hiring
- Design and tune context-engineering systems for production agents, covering retrieval pipelines, summarization and compression, working and persistent memory, and context-window management to deliver relevant information throughout multi-step workflows
- Architect the runtime harness powering production agents, including the core agent loop, tool orchestration, state management, planning and stopping logic, error recovery, and multi-agent routing, delegation, and shared-state coordination
- Define and build the action surface for production agents by integrating internal APIs, databases, ticketing systems, and program-specific platforms through custom MCP servers with model-legible, safely scoped tools. Deliver multiple production agents and MCP services across programs using Pydantic AI and LangChain/LangGraph
- Build evaluation and observability infrastructure for production agents, including evals, tracing, logging, and monitoring. Champion evals-driven development and first-class observability from project inception
- Architect guardrails, security controls, and access management for production agents, including input and output filtering, prompt-injection defenses, permissioning, and scoped identity and authorization. Enforce each user's existing access rights and address security requirements that often determine deployment approval

Lockheed Martin Corporation

Oct 2022 - Apr 2026

Orlando, FL

Staff AI/ML Engineer | Technical Lead

- Designed and delivered the organization's first production AI agents and custom MCP services across classified and unclassified programs using Pydantic AI and LangChain/LangGraph, laying the runtime, evaluation, and guardrail foundations for the current agentic platform
- Served as technical lead across AI/ML and data engineering efforts, guiding designs, reviews, and delivery as systems moved from prototype into production
- Architected a shared MLOps framework supporting several programs across classified and unclassified environments, spanning training, deployment, reproducibility, and data lineage. Cut production deployment time by roughly 85%, from about two weeks to under two days, and standardized releases across 10+ models and pipelines. Supported MLOps and DataOps teams as a subject-matter expert through technical guidance and design reviews
- Diagnosed training bottlenecks in PyTorch-based deep learning models for anomaly detection using Torch Profiler and NVIDIA Nsight. Refactored the data loader and rearchitected key layers, cutting training time by over 94% on average, from at least 24 hours to under 90 minutes
- Built distributed data pipelines with Apache Airflow and Dask using functional programming patterns, processing 100–200 GB per run across 15+ data sources and tens of terabytes over time while roughly tripling throughput for ML and analytics workflows
- Spearheaded the architecture of an Apache Arrow-native data lakehouse spanning Amazon S3 and PostgreSQL, supporting 50+ TB of time-series data and metadata with sub-second analytical queries. Integrated an ML feature store serving 10+ production models
- Redesigned and normalized a core PostgreSQL schema, cutting query response times by more than 95%, from about 10 seconds to under 500 milliseconds, and tripling supported query and transaction throughput
- Deployed and operated a secure Kubernetes cluster in AWS GovCloud supporting 20+ ML and data visualization workloads with automated scaling and 99.9% application availability in a compliance-controlled environment

Applied Theta Inc

Orlando, FL

Feb 2016 - Oct 2022

ML Engineer | Software Engineer | Data Engineer | Owner

Owner-operated ML R&D practice (sister brand of Zia Technology Group) focused on ML-based algorithmic trading research and development, with selective external ML consulting.

- Founded and directed an independent ML R&D practice, defining the end-to-end technology strategy, roadmap, and agile delivery process across data, modeling, and algorithmic trading systems
- Engineered three open-source Python frameworks spanning data preparation through training, evaluation, and serving: Photon (deep learning lifecycle), Maxwell (market-data preprocessing and ML pipelines), and Dyson (financial modeling and algorithms)
- Designed deep learning and probabilistic models (CNNs, RNNs, Transformers, Bayesian neural networks, deep ensembles) powering algorithmic trading systems across equities, options, futures, and crypto markets
- Built GPU-accelerated (CUDA) data pipelines ingesting 20+ years of trade and order-book market data at multiple time resolutions, processing hundreds of millions of data points with Dask, RAPIDS, and Apache Arrow/Parquet
- Standardized research and production codebases on TensorFlow, PyTorch, Keras, NumPy, and Pandas, eliminating duplicated implementations and streamlining the path from experimentation to live trading
- Delivered external ML consulting, including a transformer- and autoencoder-based prediction system that improved a sports-analytics client's predictive accuracy by more than 65%

Zia Technology Group Inc

Orlando, FL

Mar 2010 - Aug 2020

Data Engineer | Software Engineer | IT Consultant | Owner

External data and software consulting brand of the same owner-operator practice as Applied Theta, working with client senior leadership to shape data-management technology strategies as consultant, architect, and lead developer for businesses across industries.

- Architected data integrations and ETL pipelines connecting 25+ business applications and data sources, including MS SQL Server, Salesforce API, and QuickBooks Enterprise. These systems automated reporting workflows and eliminated hundreds of hours of manual data handling annually
- Deployed and administered 40+ SQL/NoSQL databases with custom schema designs supporting client applications, data warehouses, and reporting systems
- Delivered full-stack web applications and reporting systems, including a client web portal that unified four backend systems in real time via Node.js and WebSockets for hundreds of business users
- Designed and implemented highly available cloud and hybrid co-lo/cloud virtual infrastructures for client environments, sustaining 99.9% uptime

Thinksys Incorporated

Orlando, FL

Jan 2007 - Mar 2010

Software Developer | DBA | IT Consultant | Owner

Owner-operated consulting practice collaborating with client senior leadership to align technology with business needs while serving as consultant, DBA, and lead developer on critical IT infrastructure, data-integration, and web development projects.

- Administered 30+ MS SQL Server instances and databases powering custom and third-party business applications across client environments
- Designed and developed complex data integrations connecting multiple business applications and back-end data sources across client environments
- Developed web-based interfaces providing business users self-service access to back-end data sources, reducing manual data requests
- Led a team of 4 network engineers and developers delivering client infrastructure and software projects

Network Engineer | Software Developer | DBA | Owner

Full-stack IT consulting practice supporting networks, servers, and applications for small and mid-sized businesses across a range of industries.

- Led a team of 15 technicians and network engineers supporting 200+ client networks and applications, improving on-time delivery and service reliability across engagements
- Deployed and supported hundreds of client networks, servers, and MS SQL Server database environments for custom and third-party applications
- Built dynamic, database-driven websites and web applications for small-business clients

SELECTED PROJECTS & OPEN SOURCE

Photon – ML Framework github.com/sequenzia/photon · Jun 2018 - Oct 2022

Open-source ML framework extending TensorFlow, Keras, and PyTorch for deep learning on financial time series

- **Stack:** Python, TensorFlow, Keras, PyTorch, NumPy, Pandas, Scikit-learn, Apache Arrow, Apache Parquet, Distributed Training
- Native TensorFlow distributed-strategy support with shared inputs/outputs across networks for deep ensemble learning and dynamic learning-rate scheduling
- Real-time preprocessing for dataset splitting, normalization, scaling, aggregation, and time-series resampling, with custom batching, padding, and masking
- Simple interface for saving, serializing, and loading entire networks, including learned and hyperparameters
- Detailed parameter logging throughout the pipeline for interpretability and reproducibility

Maxwell – ML Data Pipelines github.com/sequenzia/maxwell · Jun 2018 - Oct 2022

GPU-accelerated pipelines processing 20+ years of market data with Dask, RAPIDS, and Apache Arrow

- **Stack:** Python, Pandas, NumPy, Numba, RAPIDS, Dask, Apache Arrow, Apache Parquet, Scikit-learn, NVIDIA CUDA
- Automated labeling of time-series data with algorithmic trading attributes such as stop losses, long/short price targets, ATR, and VWAP
- Domain-driven feature engineering from technical indicators and statistical markers, with automated binning, one-hot encoding, scaling, and normalization tuned for market microstructure data
- ML pipeline orchestration tools managing data flow across modeling and algorithmic-trading workflows with fluid feature extraction
- GPU-accelerated (CUDA) processing with Dask and RAPIDS for hundreds of millions of data points

Dyson – Modeling & Algorithm R&D github.com/sequenzia/dyson · Jun 2018 - Oct 2022

Financial market models spanning statistical methods through deep neural networks and probabilistic reasoning

- **Stack:** TensorFlow, Keras, Scikit-learn, statsmodels, CNNs, RNNs, Transformers, Bayesian Neural Networks, ARIMA
- Model families spanning CNNs, RNNs, Transformers, Temporal Convolutional Networks, Autoencoders, Bayesian neural networks, ARIMA, and structural time-series models
- Extensive exploratory data analysis (EDA), feature engineering, scaling/normalization, resampling, and principal component analysis (PCA)
- Advanced generalization and regularization techniques that lift model performance beyond the training data, specialized for financial time series

Algorithmic Trading System Feb 2016 - Oct 2022

Python/C++ libraries powering live execution and backtesting across equities, options, futures, and crypto

- **Stack:** Python, C++, NumPy, SciPy, Pandas, Numba, statsmodels

EDUCATION & PROFESSIONAL DEVELOPMENT

Autodidact and lifelong learner with 20+ years of self-directed, build-first professional development spanning networking, databases, software engineering, data engineering, machine learning, and agentic AI. This expertise is demonstrated through production systems, self-authored open-source ML frameworks, and a growing ecosystem of open-source agentic tools.

- Coursera (DeepLearning.AI): Machine Learning Engineering for Production (MLOps) — Oct 2022
- Google: TensorFlow Developer — Sep 2022