

Stephen Sequenzia

Orlando, Florida | stephen@sequenzia.com | 407.718.2010 | [linkedin.com/in/sequenzia](https://www.linkedin.com/in/sequenzia) | github.com/sequenzia | sequenzia.com

Active Secret Security Clearance

PROFESSIONAL SUMMARY

Senior Staff AI/ML Engineer & Architect with 24+ years building production ML, data, and software systems end to end, now focused on agentic AI in secure, large-scale defense environments. At Lockheed Martin, architected an enterprise MLOps framework and an Apache Arrow-native data lakehouse, cut deep learning training time by over 94% and PostgreSQL query times by over 95%, and drove team-wide adoption of multi-agent systems built with Pydantic AI, LangChain/LangGraph, and custom Model Context Protocol (MCP) servers on Kubernetes and AWS GovCloud. Author of three open-source ML frameworks (Photon, Maxwell, Dyson), with two decades as an owner-operator pairing deep hands-on engineering with technical leadership that turns cutting-edge AI into reliable, measurable business value.

TECHNICAL SKILLS

Agentic AI: Pydantic AI, LangChain, LangGraph, Model Context Protocol (MCP), multi-agent systems, tool-augmented reasoning, agent orchestration

Machine Learning & Deep Learning: PyTorch, TensorFlow, Keras, scikit-learn, Transformers & attention mechanisms, Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Bayesian neural networks, deep ensemble learning, anomaly detection, time-series modeling

MLOps & Cloud: MLOps (model lifecycle, reproducibility, data lineage), Kubernetes, Docker, Amazon Web Services (AWS) GovCloud, Amazon S3, distributed training, Graphics Processing Unit (GPU) / CUDA acceleration

Data Engineering: Apache Airflow, Dask, Apache Arrow, Apache Parquet, PostgreSQL, SQL/NoSQL, Extract-Transform-Load (ETL), data lakehouse architecture, ML feature stores

Languages: Python, C++, SQL

PROFESSIONAL EXPERIENCE

Lockheed Martin Corporation — Orlando, Florida | Oct 2022 – Present

Senior Staff AI/ML Engineer & Architect | Apr 2026 – Present

Staff AI/ML Engineer | Technical Lead | Oct 2022 – Apr 2026

- Serve on the leadership team of The Forge, guiding architecture across all of the team's projects (8 concurrent efforts) and shaping team growth through hiring, including 20+ technical interviews for new engineering hires.
- Lead technical direction of the Engineering AI Enablement team at Lockheed Martin Missiles and Fire Control (MFC), defining agentic AI and ML tooling strategy and enablement patterns reaching 100+ engineers.
- Architected an MLOps framework spanning the full ML lifecycle — training, deployment, reproducibility, and data lineage — cutting production deployment time by roughly 85% (from ~2 weeks to under 2 days) and standardizing releases across 10+ models and pipelines.
- Designed and built custom multi-step AI agents using Pydantic AI and LangChain/LangGraph; engineered bespoke Model Context Protocol (MCP) servers for real-time data integration with databases and object stores, delivering 6+ production agents and MCP services.
- Championed team-wide adoption of agentic AI, mentoring 10+ engineers on agent architecture and tool-augmented reasoning and cutting agent integration time by ~50% across 4 project teams.
- Diagnosed and eliminated training bottlenecks in PyTorch deep learning anomaly-detection models using Torch Profiler and NVIDIA Nsight; refactored the data loader and rearchitected key layers to cut training time by over 94% (from ~24 hours to under 90 minutes).
- Built distributed, functional-programming-based data pipelines with Apache Airflow and Dask, processing tens of terabytes across 15+ data sources and roughly tripling throughput for ML and analytics workflows.
- Spearheaded architecture of an Apache Arrow-native data lakehouse on Amazon S3 and PostgreSQL, managing 50+ TB of time-series data, metadata, and an ML feature store serving 10+ production models with sub-second analytical queries.

- Refactored and re-normalized a core PostgreSQL database, cutting query response times by over 95% (from ~10 seconds to under 500 ms) and enabling efficient scaling to ~3x higher workloads.
- Deployed and operate a secure Kubernetes cluster on AWS GovCloud supporting 10+ ML and data-visualization workloads, sustaining 99.9% availability with automated scaling in a compliance-controlled environment.

Applied Theta Inc — Orlando, Florida | Feb 2016 – Oct 2022

Machine Learning Engineer | Software Engineer | Data Engineer | Owner

Owner-operated ML R&D practice and sister brand of Zia Technology Group (common ownership, split for branding): Applied Theta focused on internal 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 — Photon (deep learning lifecycle), Maxwell (market-data preprocessing and ML pipelines), and Dyson (financial modeling and algorithms) — spanning data preparation through training, evaluation, and serving.
- 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, 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, Florida | Mar 2010 – Aug 2020

Data Engineer | Software Engineer | IT Consultant | Owner

External data/software consulting brand of the same owner-operator practice as Applied Theta, serving business clients across industries.

- Architected data integrations and ETL pipelines connecting 25+ business applications and data sources — including MS SQL Server, Salesforce API, and QuickBooks Enterprise — automating reporting workflows and eliminating 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 unifying four backend systems in real time via Node.js and WebSockets — serving 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, Florida | Jan 2007 – Mar 2010

Software Developer | DBA | IT Consultant | Owner

- Administered 30+ MS SQL Server instances and databases powering custom and third-party business applications 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.

Biziteks Incorporated — Orlando, Florida | Dec 2001 – Jan 2007

Network Engineer | Software Developer | DBA | Owner

- 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 machine learning framework extending TensorFlow, Keras, and PyTorch for deep learning on financial time series, supporting the full ML lifecycle from data preparation through training, evaluation, deployment, and serving.

Stack: Python, TensorFlow, Keras, PyTorch, NumPy, Pandas, Apache Arrow/Parquet, scikit-learn

- Custom object-oriented API with built-in subclassing of Keras/TensorFlow APIs and custom models, layers, optimizers, and loss functions; model/algorithm agnostic and container-native.
- Native TensorFlow distributed-strategy support, shared inputs/outputs across networks for deep ensemble learning, and dynamic learning-rate scheduling.
- Real-time preprocessing — dataset splitting, normalization, scaling, aggregation, and time-series resampling — with custom batching, padding, and masking, plus detailed parameter logging for interpretability.

Maxwell – ML Preprocessing & Pipelines | github.com/sequenzia/maxwell | Jun 2018 – Oct 2022

Custom libraries for acquisition, preparation, and storage of financial market data, including WebSockets/REST API connectors, anomaly detection across hundreds of millions of data points, and domain-specific feature engineering and labeling.

Stack: Python, Pandas, NumPy, Numba, RAPIDS, Dask, Apache Arrow/Parquet, NVIDIA CUDA

- Processed 20+ years of trading and order-book market data at multiple time resolutions from high-volume real-time market APIs.
- Automated labeling of time-series data with algorithmic trading attributes (stop losses, long/short price targets, ATR, VWAP), plus binning, one-hot encoding, scaling, and normalization.
- Accelerated data storage and retrieval using Apache Arrow in-memory columnar storage and Apache Parquet persistent on-disk storage.

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

Evolving collection of models and algorithms for financial markets — equities, options, futures, and crypto — spanning traditional statistical modeling through deep neural networks and probabilistic reasoning.

Stack: TensorFlow, Keras, scikit-learn, statsmodels, CNNs/RNNs/Transformers, Bayesian neural networks, ARIMA

- Developed sequence-based models producing predictive inferences of future asset prices, with custom optimizers and loss functions built for the unique complexities of time-series market data.
- Conducted extensive exploratory data analysis (EDA), feature engineering, scaling/normalization, resampling, and principal component analysis (PCA).

Algorithmic Trading System | Feb 2016 – Oct 2022

Python/C++ tools and libraries interfacing with market brokers to execute and backtest algorithmic trading across equities, options, futures, and crypto markets.

Stack: Python, C++, NumPy, SciPy, Pandas, Numba

- Unified object-oriented libraries powering both live execution and backtesting for rules-based and ML-driven strategies, with built-in exposure to key trading metrics.

PROFESSIONAL DEVELOPMENT & CERTIFICATIONS

Autodidact and lifelong learner: 24+ years of self-directed, build-first professional development spanning networking, databases, software engineering, data engineering, and machine learning — with expertise demonstrated through production systems and self-authored open-source ML frameworks (Photon, Maxwell, Dyson).

- Google: TensorFlow Certified Developer — September 2022
- Coursera: Machine Learning Engineering for Production (MLOps) — October 2022