Stephen Sequenzia

🖂 stephen@sequenzia.com 🛛 🖬 linkedin.com/in/sequenzia 407.718.2010

• Orlando, Florida

🖸 github.com/sequenzia

sequenzia.com

A skilled staff-level Machine Learning Engineer with a strong foundation in Software Engineering, Data Engineering, and MLOps. With hands-on experience building end-to-end machine learning pipelines, I specialize in designing scalable data architectures, deploying robust ML models, and automating workflows using MLOps best practices. Passionate about creating impactful AI solutions, I leverage my expertise in Python, cloud services, and DevOps tools to bridge the gap between data science experimentation and production deployment. My goal is to deliver reliable and efficient machine learning systems that drive business value.

EXPERIENCE

Lockheed Martin Corporation Staff Machine Learning Engineer | Technical Lead

As both a technical lead and individual contributor specializing in machine learning and data-centric systems, I have driven technical initiatives and led a small but diverse team of software, machine learning, and data engineers to deliver innovative solutions that align with best practices and business objectives. My focus has been on fostering collaboration, promoting technical excellence, and enhancing team capabilities to achieve impactful results in AI and data projects.

- Developed a robust MLOps framework with tools like Ray, Optuna, Weights & Biases, and MLflow, optimizing the entire ML lifecycle from model training to deployment. This framework significantly accelerated deployment, improved model reproducibility, and strengthened data lineage while reducing the complexity and time required for deploying to production.
- Leveraged tools such as Torch Profiler and Nvidia Nsight to diagnose and address inefficiencies in PyTorch-based deep learning anomaly detection models. By refactoring the data loader and rearchitecting key layers, achieved a reduction in training times of over 94%.
- Built highly distributed, functional programming-based data pipelines with Apache Airflow and Dask to support complex machine learning and analytics workflows, enabling efficient, scalable data processing and seamless integration of diverse data sources for enhanced throughput and reliability
- Principal architect of an Apache Arrow-native Data Lakehouse using S3 and PostgreSQL for efficient management of time series data, metadata, and a machine learning feature store, enabling seamless access, high-performance queries, and scalable analytics.
- Refactored a PostgreSQL database for improved normalization, performance, and scalability, reducing query response times by over 95%. Optimizations streamlined data access, minimized redundancy, and enabled efficient scaling to meet higher workloads.
- Deployed and managed a secure, scalable Kubernetes cluster on AWS GovCloud to support ML and data visualization platforms, ensuring high availability, automated scaling, and efficient resource management in a compliant environment.

Applied Theta Inc

Machine Learning Engineer | Software Engineer | Data Engineer | Owner

Develop analytic opportunities through a suite of frameworks and models designed to expedite development and uphold the integrity of applied AI/ML models tailored to the specific demands of financial markets. Originally built for internal purposes, many of these models and frameworks are now available as open-source resources.

- Developed the entire technology strategy, roadmap, and solutions using agile software development.
- Developed several python-based frameworks and libraries.
- Utilized Machine Learning algorithms to produce financial models and algorithmic trading systems.
- Built advanced data pipelines to handle data from various market sources and time resolutions.
- Standardized codebases to utilize leading data science frameworks like TensorFlow, Keras, NumPy, and Pandas.
- Provided Machine Learning based consulting, design and development to other companies.

Orlando, Florida Oct 2022 – Present

Orlando, Florida Feb 2016 – Oct 2022

Zia Technology Group Inc Data Engineer | Software Engineer | IT Consultant | Owner

Worked with senior leadership to shape technology strategies for data management, acting as consultant, architect, and lead developer on solution design and development using diverse technologies and agile methods.

- Designed and developed complex data integrations across multiple business applications.
- Deployed many SQL/No SQL databases and designed custom data schemas.
- Designed and deployed data pipeline architectures across multiple data sources and developed ETL processes.
- Full-stack development of several custom web-based applications and reporting systems.
- Designed and implemented highly available virtual infrastructures, fully cloud-based and hybrid co-lo/cloud.

Thinksys Incorporated

Software Developer | DBA | IT Consultant | Owner

Collaborated with senior leadership to align technology solutions with business needs, serving as consultant, DBA, and lead developer on critical IT infrastructure projects.

- Set up, deployed, and managed many MS SQL Servers and databases for both custom and 3rd party applications.
- Designed and developed complex data integrations across multiple business applications.
- Developed and deployed web-based interfaces for access to back-end data sources.
- Managed a small team of network engineers and developers.

Biziteks Incorporated

Network Engineer | Software Developer | DBA | Owner

Provided full-stack IT consulting, support, and development for small to medium-sized companies in various industries.

- Managed a team of 15 technicians and network engineers.
- Set up, deployed, and managed many MS SQL Servers and databases for both custom and 3rd party applications.
- Set up, managed, and supported hundreds of client networks and applications.
- Developed web-based applications and dynamic websites.

PROJECTS

Photon ML Framework

Photon ML is a Machine Learning framework extending TensorFlow, Keras, and PyTorch for deep learning in financial algorithms. It supports the full ML lifecycle: data preparation, model development, training, monitoring, evaluation, deployment, and serving.

Project Repo: github.com/sequenzia/photon

Technologies & Methods: Python, Neural Networks, Deep Learning, Distributed Training

Frameworks & Libraries: TensorFlow, Keras, PyTorch, NumPy, Pandas, Apache Arrow, Apache Parquet, Scikit-learn

Key Features of Photon ML:

- Streamlines the development and implementation of end-to-end Machine Learning systems.
- Custom object-oriented API with built-in subclassing of Keras and TensorFlow APIs.
- Built-in custom modules such as models, layers, optimizers, and loss functions.
- Highly customizable interface to extend built-in modules for specific algorithms/networks.
- Detailed logging and analysis of model parameters to increase interpretability and optimization.
- Works natively with TensorFlow distributed strategies.
- Real-time data preprocessing; dataset splitting, normalization, scaling, aggregation, and time series resampling.
- Custom batching, padding, and masking of data.
- Designed to be model/algorithm agnostic and to work natively with container services.
- Natively shares input and output between multiple networks to streamline deep ensemble learning.
- Simple interface for saving, serializing, and loading entire networks including learned and hyperparameters.
- Custom dynamic learning rate scheduling.

Orlando, Florida

Jan 2007 - Mar 2010

Orlando, Florida Dec 2001 - Jan 2007

June 2018 – Oct 2022

Orlando, Florida Mar 2010 - Aug 2020

Modeling/Algorithm Research & Development

An evolving collection of data models and algorithms used to model financial markets including equities, options, futures, and crypto assets. The project has grown from only traditional statistical modeling to Machine Learning based modeling which includes deep neural networks and probabilistic reasoning networks.

Project Repo: github.com/sequenzia/dyson

Technologies & Methods: Deep Learning, Neural Networks, Regression, Classification, Probabilistic Reasoning, Deep Ensemble Learning, Back-Propagation, Statistical Analysis, Financial Modeling

Frameworks & Libraries: TensorFlow, Keras, Pandas, NumPy, Seaborn, Matplotlib, Statsmodels, Scikit-Learn

Algorithms & Models: CNNs, RNNs, Transformers, Attention Mechanisms, Temporal Conv Networks, Autoencoders, Bayesian Neural Networks, ARIMA, Structural Time Series

- Developed algorithms that used sequence-based modeling to make predictive inferences of assets future prices.
- Utilized advanced generalization and regularization techniques to increase the performance of models.
- Developed custom optimizers and loss functions to address the unique complexities of time series market data.
- Conducted extensive exploratory data analysis (EDA), feature engineering, data scaling/normalization, resampling, and principal component analysis (PCA).

Machine Learning Preprocessing & Pipelines

June 2018 – Oct 2022

Developed a set of custom libraries to handle the unique characteristics of acquisition, preparation, and storage of financial market data. These libraries include WebSockets/RESTful API data connectors to access data, detect anomalies in hundreds of millions of data points, then cleanse and pre-process to provide high-integrity data modeling. Also included is a set of tools for domain-specific feature engineering and labeling of financial market data.

Project Repo: github.com/sequenzia/maxwell

Technologies & Methods: Data Preprocessing and Cleaning, Feature Engineering, Machine Learning Pipelines, Data Labeling, WebSockets/RESTful APIs, Nvidia CUDA GPUs, Python Data Structures, and Storage

Frameworks: Pandas, NumPy, Numba, Rapids AI, DASK, Apache Arrow, Apache Parquet, Scikit-learn, Seaborn, Matplotlib

- Utilized domain knowledge to engineer features based on technical indicators and statistical markers.
- Processed over 20 years of both trading and book market data with multiple time resolutions.
- Increased data storage and retrieval speed and efficiency, utilizing Apache Arrow for in-memory columnar storage and Apache Parquet for persistent on-disk storage.
- Produced a wealth of useful tools for managing the flow of data for various modeling and algorithmic trading, orchestrating Machine Learning pipelines, providing a high level of fluidity to feature extraction, and helping to increase the integrity of the data.
- Designed custom procedures to label time series data with algorithmic trading attributes such as stop losses, long/short price targets, ATR, and VWAP values.
- Processed and stored data from high-volume real-time market data APIs.
- Automated data labeling, binning, one-hot encoding, scaling, and normalization of time series market data.

Algorithmic Trading System

Feb 2016 – Oct 2022

Developed a set of tools and libraries designed to interface with market brokers and execute algorithmic trading. The algorithmic trading systems execute both buy and sell orders in equities, options, futures, and crypto markets. The tools also support backtesting of trading algorithms, applying predictive models to historical data to evaluate their accuracy.

Technologies & Methods: Python, C++, Algorithmic Trading, Financial Models

Frameworks: NumPy, SciPy, Pandas, Numba, Seaborn, Matplotlib, Statsmodels

- Developed utilizing object-oriented programming in Python and C++ for increased performance.
- Designed a set of libraries that work with both rules-based and Machine Learning based models and algorithms.
- These same libraries can be used for both live executions of algorithmic trading systems and backtesting.
- Produced a highly functional and efficient algorithmic trading and backtesting system; able to trade various assets in various markets and provide increased exposure to key trading metrics.

June 2018 – Oct 2022

Machine Learning Consulting & Development

Collaborated with the company owner, a data professional, to transition from traditional statistical analysis to machine learning models, enhancing inference quality and prediction accuracy for future sporting events—the company's core focus. Leveraged attention-based transformers, achieving over a 65% improvements in predictive accuracy.

Technologies & Methods: Deep Learning, Neural Networks, TensorFlow, Predictive Modeling, Transformers, Autoencoders, Statistical Analysis

- Developed a custom pipeline to extract Machine Learning focused features from a dataset with many manual attributes built into the data.
- Designed and developed a TensorFlow/Keras based variational autoencoder neural network to reduce the dimensionality of the data.
- Designed and developed custom TensorFlow/Keras based models, algorithms, and an attention-based transformer network to make inferences/predictions in a highly stochastic environment.

Data Warehousing & Integrations

Collaborated with the CFO and department heads of a national court reporting firm to resolve data integration issues impacting workflow and efficiency. As architect and lead developer, designed solutions to unify data silos, improving integration and automating key accounting processes.

Technologies & Methods: MS SQL Server, ETL, Data Warehousing, Batch Processing, MS Integration Services

- Budget management, training, and supervisory responsibilities (employees and consultants).
- Built framework to batch process large data sets and store them in a cloud-based data warehouse.
- Reverse engineered the database structure of a closed system and developed internal schemes and documentation.
- Delivered a product that provided real-time data integrations across business applications and the backend storage systems (MSSQL Server, Salesforce API, and QuickBooks Enterprise).

Data Modeling, Analytics & Reporting

Collaborated with the CFO and department heads of a national court reporting firm to address information needs during rapid growth. Developed revenue prediction models, key performance dashboards, and a reporting system for consolidated financials, supporting executives and external stakeholders.

Technologies & Methods: Model Development, Data Analytics, Reporting, Data Visualization

- Developed models to predict future revenues during high growth in many new geographic areas. Utilized multiple linear regression and other predictive modeling strategies to accurately forecast future revenues.
- Took on an expanded role, diving deep into the business side of the project.
- Blended business knowledge and revenue models with a very strong understanding of multi-dimensional data analysis to develop models and an analytical dashboard to support corporate acquisition strategies.

Web Portal & SQL Databases

Developed and supported a custom web app that was used as a hub for connecting four backend systems across multiple lines of business.

Technologies & Methods: SQL Databases, Search Algorithms, PHP, Node.js, XML, WebSockets

- Designed and developed the entire technology stack including SQL databases, and PHP/JavaScript interfaces.
- Integrated real-time data from external databases utilizing Node.js and WebSockets.
- Integrated data from QuickBooks Enterprise utilizing a custom-developed XML/PHP framework.
- Designed several search algorithms, including binary search and fuzzy match search.

CERTIFICATIONS

Google: TensorFlow Certified Developer

Coursera: Machine Learning Engineering for Production (MLOps)

Sept 2021 – Dec 2021

September 2022

October 2022

2016 – 2019 tion issues

2014 – 2019

2014 - 2019