



Pavlo H

Expert Senior ML/AI/MLOps Engineer

SUMMARY

Experienced IT professional with more than ten years of expertise in Cloud Systems, DevOps, Machine Learning and MLOps practices. Pavlo specializes in building scalable and reliable ML pipelines, optimizing models for real-world use, and managing full lifecycle of machine learning projects. Pavlo's technical skills include TensorFlow, PyTorch, Kubernetes, Kubeflow, Terraform, Python, as well as cloud platforms and services such as AWS SageMaker and Azure Cognitive Services.

TECHNICAL SKILLS

Main Technical Skills	TensorFlow (5 yr.), PyTorch (5 yr.), Scikit-learn, Kubeflow, Mlflow
Programming Languages	Java, Python
Java Frameworks	Apache Spark, Hibernate
Scala Frameworks	Apache Spark
AI & Machine Learning	BERT, COLBERT, Kubeflow, Mlflow, PyTorch (5 yr.), Scikit-learn, Tensorboard, TensorFlow (5 yr.), Triton, Vertex AI, YOLOv5, YOLOv8
PHP Libraries and Tools	Composer
Python Frameworks	Django, Flask
JavaScript Libraries and Tools	Gulp.js, Webpack
Python Libraries and Tools	Pandas, pytest, PyTorch (5 yr.), Scikit-learn, TensorFlow (5 yr.)
JavaScript Frameworks	Vue.js
Data Analysis and Visualization Technologies	Apache Airflow, Apache Hive, Apache Spark, Pandas, Teradata
Databases & Management Systems / ORM	Apache Hadoop, Apache Hive, Apache Spark, AWS DynamoDB, Hibernate, Microsoft Azure SQL Server, MongoDB, PostgreSQL, Redis

UI Frameworks, Libraries, and Browsers	Bootstrap 4
Azure Cloud Services	AKS, Azure Blob Storage, Azure Cognito, Microsoft Azure SQL Server
Amazon Web Services	AWS DynamoDB, AWS EC2, AWS Lambda, AWS S3, AWS SQS
Google Cloud Platform	GCP Dataflow
Deployment, CI/CD & Administration	Ansible, Jenkins, Pulumi
QA, Test Automation, Security	Apache Maven, pytest, Unit Testing
Web/App Servers, Middleware	Apache Tomcat
Scripting and Command Line Interfaces	Bash
Third Party Tools / IDEs / SDK / Services	CMake, GCC, GDB, valgrind
Virtualization, Containers and Orchestration	Docker, Kubernetes, Terraform
Message/Queue/Task Brokers	Kafka, RabbitMQ

PROFESSIONAL EXPERIENCE

LEAD ML ENGINEER, DEC 2023 – NOW

As a Lead ML Engineer, my role has been instrumental in defining the technical direction and driving the development of innovative solutions for a media startup. In this dynamic and fast-paced environment, I focused on building a strong, collaborative team capable of tackling complex challenges, particularly in designing and implementing advanced ML pipelines.

In addition to team lead responsibilities, I have been deeply involved in crafting scalable, production-ready ML infrastructure, ensuring efficient data flow, model training, and real-time inference. My work included introducing best practices in pipeline automation, optimizing workflows for latency and reliability, and integrating modern tools like Kubernetes, Kafka, and Elasticsearch to enhance overall system performance. By aligning technical efforts with business goals, I ensured the delivery of solutions that not only addressed immediate client needs but also established a robust foundation for future scalability and innovation.

Responsibilities:

- Design and implement legacy client solutions (RTS/Fraud detection)
- Building Inference API's for ML outputs.
- Degradation issue detection and further optimization.
- Perform Katib model custom hyperparamethrisation experiments.
- Writing KFP (Kubeflow) pipelines for data model embeddings merger.
- Elaborate with a team to resolve critical tasks and production issues.
- Deployment stages control and phase of model testing (A/B testing).



- Establish Monitoring and Alerting for client infrastructure
- Documenting the latency issues discovery and creating PoC action items
- Analyze applied algorithm/suggest better algorithm for model training.
- Managing data flow and caching for ML training and inference with Kafka and Redis.
- Storing and retrieving large ML datasets (features) with Elasticsearch.

Achievements:

- Established monitoring for ML models to ensure performance.
- Developed a service with ML for personalized mobile notifications.
- Resolved complex issues in ML models, enhancing reliability. Managed ML jobs and data processing tasks efficiently.
- Proposed better RTS-model solution for latency model optimization (YOLOV8)
- Conducted internal RAG workshop.

Completed Goals:

- Successfully implemented RTS and fraud detection systems, achieving a measurable reduction in false-positive rates and improving decision-making latency.
- Deployed scalable inference APIs that support high-throughput, low-latency ML outputs in production environments.
- Designed and tested a custom pipeline for data embedding merges, reducing training time by 20%.
- Automated monitoring and alerting systems using Prometheus and Grafana, ensuring minimal downtime and faster incident response times.
- Documented latency bottlenecks and delivered actionable PoC solutions that led to a 25% reduction in pipeline execution time.
- Improved ML training efficiency by designing better algorithms and introducing optimized caching mechanisms.
- Delivered comprehensive documentation, streamlining onboarding for new engineers and enhancing collaboration across teams.

Key Technologies: AWS (Sagemaker, Bedrock), Kubeflow, Tensor RT, Prometheus, Grafana, Azure Cognitive Service, Python (FastAPI, pandas, matplotlib, pytorch)

AI ENGINEER, ML Engineer, ReFaceApp – NOV 2021 – DEC 2023

Reface is a dynamic startup specializing in AI and IoT solutions for smart artificial vision systems. I played a key role in building the embedded systems team, delivering robust software for microcontroller-based applications. Our work involved close collaboration with hardware teams to design, develop, and implement software for a variety of IoT devices, with a strong emphasis on creating smart, efficient, and energy-saving systems.

Team Size: 5 ML, 2 BE, 1 FE, 2 QA

Project Roles: ML/MLOps Engineer Responsibilities:

- Leading the Embedded Systems team, ensuring efficient and timely delivery of projects.
- Designing, developing, and testing software embedded in devices and systems.
- Integrating and validating new product designs.
- Enhance signal processing algorithms.
- Design and implement BERT transformers for natural language processing tasks
- Supporting software QA and optimize I/O performance.
- Develop ML pipelines and monitoring.
- Managing C, C++, and Python development for embedded systems.
- Working closely with the hardware team to ensure software and hardware integration.



- Providing post-production support for embedded systems.
- Engaging in code and design reviews to maintain our high development standards.
- Collaborating with cross-functional teams to define, design, and ship new features.

Achievements:

- Successfully developed and launched a line of IoT products related to home automation that were praised for their reliability and user-friendly interface, resulting in a 30% increase in company sales.
- Designed and implemented BERT transformers for natural language processing tasks, including tokenization, embedding generation, model experiments and benchmarking.
- Created an IoT communication bridge between emCOM with LLM embedding model.
- Optimized slow model that showed degradation and performance issues.
- Implemented a lightweight MQTT protocol to enable IoT devices to communicate efficiently with backend servers, enhancing real-time remote device management and monitoring.
- Improved the operational efficiency of devices by optimizing embedded code and implementing energy-saving features, ensuring longer device life and customer satisfaction.
- Guided a team through the full product development lifecycle, from initial concept to production and support.

Key Technologies: C, Vertex AI, Tensorflow, C++, EfficientNet, Python, MQTT, Embedded Systems, Microcontrollers (e.g., Arduino), IoT, Real-Time Systems, Multithreading, GCC, GDB, Valgrind tool, cMake, Network/Hardware protocols.

STAFF MLOPS ENGINEER, LUXOFT (DXC) – JUN 2020 - NOV 2021

As a contractor, I contributed to a renowned global logistics company in the US, working as a Senior Engineer. My primary responsibility was to deliver high-quality, consistent products to stakeholders while maintaining and optimizing a complex system of ML and ETL pipelines.

These pipelines were critical for analyzing client preferences and forecasting seasonal goods sales.

In this role, I played a key part in enhancing the functionality and efficiency of the system by optimizing data workflows across tools like MLflow, SageMaker, Airflow, Kafka, and proprietary internal frameworks. My efforts ensured the seamless integration and performance of these components, enabling more accurate and reliable business insights.

Team Size: 3 MLOps, 1 BE, 1FE, 2 DS

Project Roles: MLOps Engineer/Staff MLOps Engineer

Responsibilities:

- Improving ML pipeline DAG's tasks for datasets augmentation (MLflow + Airflow).
- Optimizing and fixing general failing/timing/sync issues (cluster resources profiling).
- Development and support of existing ML Pipelines.
- AWS stack (S3, SQS, EC2, Glue, Sagemaker, Bedrock).
- Analyze for possible model degradation. Perform experiments in Kubeflow along with KFP
- Distributing code to IsF Framework (internal infrastructure).
- Integrating third party API's.
- Estimating model performance by its metrics (F1 Score, AUC, type 1 and type 2 errors).
- Model drift and bias monitoring.
- Establish Monitoring and Alerting for client infrastructure.



Achievements:

- Implemented C extensions for ISF framework (CV core).
- Legacy support and feature requests developing.
- Speed up model inference for Data Platform high availability and data consistency.
- Balancing and optimizing k8 clusters for better model latency.
- Database optimization and reliability check.
- Writing bash/terraform scripts.
- Writing component tests and conduct A/B testing.
- Work close with Data Science team to obtain common goals.

Accomplishments:

- Developed functionality for data processing internal framework (Python and Java code bases).
- Fixed ETL pipelines occurrences/tasks stuck (every 2 weeks for 12h - notification came from Pagerduty).
- Wrote DDLs for hive/teradata.
- Wrote infrastructure scripts and code base.
- Conducted two successful interview for MLOps team hiring process.
- Optimized code base (performance and app reliability).
- Wrote documentation for AWS internal infra.
- Created documentation for framework.
- AWS cost optimizations.

Key Technologies: AWS (Sagemaker, Bedrock, Q), MLFlow, Tensor RT, TensorFlow Serving, Prometheus, Azure Kubernetes Service, Python (FastAPI, pandas, matplotlib, pytorch, numpy)

SOFTWARE ENGINEER / DEVOPS ENGINEER, FDNA INC – MAY 2020 - DEC 2020

Gained foundational experience in Machine Learning (ML) while working on the Gestalt Match project for precision medicine, which leverages AI algorithms. My primary responsibilities included supporting and building CI/CD pipelines for the project, maintaining and enhancing existing systems, and addressing critical bug fixes.

Contributed to the project by distributing new library components using CondaForge to build recipes, serving as the raw material for advanced solutions. Worked on complex tasks involving

PyTorch and TensorFlow to process and distribute datasets for the project's end-to-end AI-driven features, further solidifying my expertise in ML technologies.

Responsibilities:

- Integrating Payment System.
- Model experiments degradation monitoring.
- Improve unittest/E2E test coverage.
- Triggering webhooks for user actions.
- Sending emails according to the schedule.
- Notification system implementation.
- Ads targeting system maintaining.
- Code reviews and multi-staging approvals.
- Maintain Jenkins infra deployment pipeline.
- Distributed import module for CSV/XLS files.
- Developing accountant data system.



Accomplishments:

- Developed functionality for internal CI/CD integration tool.
- Auditing applications (measuring performance, code base quality, security etc.).
- Optimization of code base.
- Garbage collection optimization during data pipeline ingestion.
- Written bunch scripts for infrastructure.
- DAG data platform ingestion speed improvement.
- I gained experience in Big Data processing.

Key Technologies: Snowflake, DBT, Airflow, AWS (Lambda, Redshift, S3, Glue), Python (FastAPI, neo4j, pandas, numpy, multithreading, asyncio), Triton Inference Server

PYTHON DEVELOPER, LITSLINK – 2019 - MAY 2020

Worked on diverse projects spanning industries such as e-commerce, insurance, and investment, delivering customized solutions using Django REST Framework and AWS. Successfully initiated and developed projects from the ground up, ensuring they were thoroughly tested, well-maintained, and aligned with client requirements.

Designed and implemented robust APIs from scratch, incorporating complex business logic to enhance security, reliability, and scalability. Additionally, acquired significant experience in optimizing and refactoring legacy codebases by upgrading outdated Python libraries and extensions, resulting in improved performance, maintainability, and alignment with modern development standards.

Project Roles: Python Engineer

Responsibilities:

- Developing RESTFUL API's.
- Billing system integration.
- Working with Vue.js framework.
- Align codebase to CBS style code.
- Working with queues and caching (Kafka, Redis).
- Elasticsearch implementation as database.

Completed goals:

- Implementing monitoring systems.
- Implemented mobile multi-purpose notification integration service.
- Bug fixing with challenges.
- Worked with queues and tasks jobs.

Key Technologies: Django REST Framework, FastAPI, Python (asyncio, psycopg, numpy, pandas, aiohttp, multiprocessing), SQL

SOFTWARE DEVELOPER, TRAFFIC LABEL – 2018 - DEC 2019

Joined a dynamic and ambitious team as a PHP Developer, contributing to an advertisement project with direct communication and requirements provided by the client. Collaborated with the team to build a platform designed to accumulate traffic and create a dashboard for monitoring and measuring performance.

Played a key role in addressing challenges associated with high-load systems, focusing on resource optimization, scalability improvements, and ensuring system stability. Additionally, took part in DevOps activities, including working with VMware vSphere to manage virtualized environments and ensure reliable infrastructure performance under demanding conditions.



Team Size: 1 Team Lead, 2 BE, 1 FE, 1 QA Project Roles: Software Developer (Python)

Responsibilities:

- Integrating 'Toffee' system to client billing interface.
- Provide documentation.
- Pointing aggregating data from client microservices.
- Optimize queries.
- Penetration testing (Vulnerability monitoring).
- Bash scripting.
- Support high loaded system via vSphere and Jenkins.

Completed goals:

- Developed a brand-new website from scratch.
- Proposed motivational approach for application scaling.
- Hard testing (probation, simulating random attacks).
- Conducted reliability and security optimizations.
- Maintained and supported availability for a Toffee application.

Key Technologies: Composer, Laravel 5, VueJS, Jenkins, Terraform, Docker

PHP DEVELOPER, MWDN – 2016 - SEP 2018

Worked at MWDN, a company specializing in e-commerce solutions for media and product-based businesses. Focused on developing and customizing buying platforms and blogs to meet client needs.

Primarily worked with popular CMS platforms such as WordPress, OpenCart, and Magento, tailoring solutions to enhance functionality, user experience, and scalability.

Project Roles: Software Developer, LAMP stack

Responsibilities:

- Developing responsive e-commerce solutions for clients.
- Developing plugins for WordPress Engine.
- User experience testing and feedback.
- Interact with UI/UX designers.
- Site score improvements (SEO metrics).
- Support and bug fixing critical issues.

Completed goals:

- Supporting and developing feature requests for client-side sites (Sparta, Maison projects).
- Hosting support (troubleshooting).
- Bug fixing and developing e-commerce business features (payment integrations, ads targeting).

Key Technologies: PHPbb, Wordpress, jQuery, JS (ES6), gulp, WooCommerce, Magento

INDEPENDENT CONTRACTOR – 2015 - NOV 2019

Gained extensive experience as an independent contractor, completing diverse tasks for clients across various industries.

Developed over 20 websites from scratch using different CMS platforms, demonstrating expertise in building functional, user-friendly, and scalable solutions. Transitioned to Python



development in 2018, shifting from PHP to focus on more advanced technologies and frameworks.

Contractor engagements provided invaluable experience, sharpening my skills in problem-solving, adapting to client requirements, and delivering high-quality software solutions. This journey significantly contributed to my growth as a proficient and versatile Software Developer.

Project Roles: Full-stack Developer (Python, PHP, JS, GoLang)

Responsibilities:

- Complete different tasks from clients, commonly implementing new functionality to web applications / websites.
- Bug Fixes and urgent patches
- Restful API's
- Integrating third party apps
- Developing PoC from client side
- Supporting compound infrastructures
- UI design and styling

Completed goals:

- Developed bunch of sites from scratch.
- Developed RESTFUL applications for mobile.
- Maintained and worked on old code base applications (legacy support).
- Made basic mockups for SPA sites.
- Gained invaluable experience for future carrier ladder.

Key Technologies: Python, JS, PHP, Java, Golang

EDUCATION

- **Mykhailo Ostrogradsky University**, Automation Engineer, Master's Degree, 2019

CERTIFICATION

- **Microsoft Certified: Azure Solutions Architect Expert**
- **Microsoft Certified: Azure AI Engineer Associate**
- **AWS Certified Machine Learning Engineer – Associate**
- **Google Cloud Professional Cloud Developer**
- **Microsoft Certified: Azure Administrator Associate**
- **Microsoft Certified: Azure AI Fundamentals**
- **Microsoft Azure Security**
- **Microsoft Entra ID for Administrators**
- **Machine Learning in Production**
- **AI-Assisted Engineering 2.0**
- **MLOps Essentials: Model Development and Integration**
- **CS50 Harvard**
- **Python Data Structures and Algorithms**
- **Python: Programming Efficiently**
- **Python Developer**
- **Essential Java**
- **Introduction Java**
- **Digital Marketing**

