

[+48 732 839 674](tel:+48732839674)

MAKSYM VOITKO

max.voitko@gmail.com
[linkedin.com/in/mvoitko](https://www.linkedin.com/in/mvoitko)

Languages and Technologies

Python, Golang, SQL, GCP, AWS, k8s, Kubeflow, Terraform, Helm, MLOps, LLMOps, CUDA, Ray, Docker, Kafka, Pytorch, Pandas, ETL, Airflow, OpenCV

Employment

Honeycomb Software

Jul 2023 – Present

Head of Data and AI

- Own company-wide AI strategy and P&L, prioritising 6-month revenue & cost-saving bets with C-suite.
- Architected distributed embedding & inference clusters (ONNX + Triton) on AWS EKS; auto-scaling and vector-DB sharding cut compute costs by 45 %.
- Formulated and executed a comprehensive data engineering and AI strategy, defining milestones and KPIs to increase data accuracy by 25% and reduce model training time by 30%.
- Centralized monitoring and alerting with Victoria Metrics and Grafana, improving incident response time by 75%.
- Founded and directed the Data and AI department, carefully selecting and nurturing a team of 12 skilled engineers, which boosted company revenue by 20%.
- Introduced quarterly AI/ML roadmap planning, aligning technical initiatives with business OKRs and improving stakeholder engagement.
- Instituted mentorship tracks to develop junior engineers into senior contributors, resulting in 3 promotions within the first year.
- Coordinated delivery of 8+ end-to-end data/ML products, working with engineering, product, and design teams to meet aggressive deadlines without scope compromise.
- Streamlined product development process by leveraging RICE framework to identify and eliminate 50% of non-essential features; accelerated time-to-market for critical products and increased overall team productivity by 40%.
- Introduced feature estimation with PERT, resulting in on-time program delivery.
- Orchestrated the development and integration of a sophisticated multimodal LLM agent for project management, enhancing lead generation by 60%.
- Created MLOps platform based on AWS Sagemaker and RayTune with Terraform, Bash, and Python, covering all the pipeline stages from experimentation to monitoring.
- Improved the inference performance of the CV models with GStreamer by 40%.
- Implemented data pipeline for processing IoT sensor data and generating NFTs.
- Developed the data visualization tool, which increased client onboarding speed by 75%.
- Defined product metrics for experiments and production experience - dropped troubleshooting time by 20%.
- Shipped text2sql and RAG solution from POC to production service.

Gr8 Tech

Aug 2021 – Jun 2023

R&D Team Lead

- Led a cross-functional team of 10 ML and Software engineers with projects CLTV prediction, churn prediction, platform infrastructure load forecasting, and recommendation system.
- Designed and developed a Machine Learning Recommendation Platform serving tens of models to more than 2 million MAUs that increased user engagement and retention.
- Built MLOps infrastructure with K8s, Kubeflow, Airflow, and DVC, allowing the creation of 6 recommendation model types after hundreds of A/B tests and experiments.

-
- Architected and implemented an Airflow DAG Factory enabling Data Scientists and Analysts to add new DAGs as a YAML configuration with little to no code, which reduced data preparation delivery time by 75%.
 - Introduced engineering interview flow, which allowed for raising the bar of the candidates with a decreased interview loop.

SoftConstruct

Mar 2020 – Aug 2021

AJNA - an AI-powered augmented live video streaming tool that recognizes the game flow of sporting events, provides advanced player tracking in live mode, and collects more technical data to create new unique markets to bet on.

Team Lead

Jan 2021 – Aug 2021

-
- Guided a team of 3 people in implementing a broadcast video streaming service from a bespoke camera based on Nvidia Jetson; successfully progressed from a POC to the first production release for a live broadcast of soccer matches.
 - Implemented smooth PTZ frame transition algorithm for broadcast.
 - Engineered real-time panorama stitching of dual 4K frames using OpenCV on both CPU and GPU with CUDA, achieving a 40% reduction in processing time and enhancing overall performance.

Software Engineer

Mar 2020 – Jan 2021

-
- Moved the main post-processing video Computer Vision pipeline to the Google Cloud services.
 - Designed and developed the first customer-facing statistics service.
 - Improved the CV video processing pipeline component performance by 54% (99th percentile).

glomex GmbH

Jul 2017 – Mar 2020

glomex provides a one-stop solution for video publishing — from transcoding a video file, metadata processing, and providing a player for all platforms to monetization.

Software Engineer / Data Engineer

Aug 2018 – Mar 2020

-
- Migrated Data infrastructure from Data Warehouse to Data Lake.
 - Implemented a billing API service that optimized one week of Finance team calculations.
 - Developed a state-of-the-art NLP articles category classifier, reducing content categorization time by 70% and enhancing user experience by providing more accurate search results.
 - Implemented real-time data ingest flow with AWS Kinesis services, enabling the processing of 10,000+ data points per second and reducing data processing latency by 60%.
 - Improved data infrastructure, which reduced service latency 3 times from 9s to under 2s.
 - Automated reports generation for the Sales team and top management, which took 3 eng. hours/day.

Test Automation Engineer

Jul 2017 – Aug 2018

-
- Refactored test automation framework, which saved 25% of the maintenance time.
 - Automated e2e smoke test run, which saved 0.33 engineering hours after each deployment.
 - Trained and mentored 2 new members of the team.

Quality Assurance Engineer

Ignite at Waze, Alphabet

Jan 2017 – Jul 2017

Waze Carpool

-
- Led quality assurance activities on the Waze Carpool Rider app.
 - Cut regression cycle time twice by switching to checklist test documentation.
 - Mentored new members of the team.

Quality Assurance Engineer

ex-Cogniance / Star

Sep 2015 – Jan 2017

Home management web app

-
- Cut regression cycle time twice by introducing a risk-based testing strategy
 - Tested Back End performance with JMeter, which led to choosing larger instances to meet requirements.

Infotainment solution for connected cars

-
- Established and implemented comprehensive QA processes and procedures, leading to a 40% reduction in software bugs and a 30% reduction in time-to-market by 30%.

Education

Kyiv, Ukraine

Igor Sikorsky Kyiv Polytechnic Institute

Sep 2017 – Feb 2019

M.S.E. in Computer Engineering. Master's thesis: CNNs for network users' biometric identification.