Mislav Jurić

Curriculum Vitae

PERSONAL INFORMATION

Website	http://www.mislavjuric.com
LinkedIn	https://www.linkedin.com/in/mislav-juric
Email	mislav.juric01@gmail.com

WORK EXPERIENCE

Machine Learning Engineer and Consultant (remote) October 2023 - present
Self-employed

Key points:

- Self-employed ML engineer and consultant solving real-world business problems through the right ML approaches: whether classical algorithms or modern deep learning.
- Operates across the full spectrum: from low-level ML engineering to AI strategy and business alignment.
- Built and fine-tuned LLM-powered systems across multiple client engagements, including an LLM-powered game with prompt engineering and safety systems (GoodAI - AI People, see demo) and enterprise chatbots on Google's Gemini Enterprise platform.
- At Newfire, optimized inference time of BERT-based models, built new product features and maintained production NLP pipelines.
- Experience spans NLP, computer vision, traditional ML, and speech systems (TTS/ASR).
- Capable of handling full pipeline (MLOps, backend, deployment) when needed.
- Delivered lectures at AI Center Lipik (AI Centar Lipik), educating future AI developers.

Job description:
As a self-employed machine learning engineer and consultant, I solve business problems by selecting and implementing the right ML approaches for each use case. My work spans from hands-on ML engineering (training models and writing production code) to strategic consulting where I assess AI feasibility, define roadmaps, and align technical solutions with business objectives. On the LLM and GenAI side, I contributed to AI People at GoodAI (an LLM-powered game) by fine-tuning and deploying LLMs locally, engineering prompts and safety systems, implementing TTS/ASR capabilities, and maintaining cloud infrastructure. In a separate engagement, I built enterprise chatbots on Google's Gemini Enterprise platform, working with ADK (Agent Development Kit) and Google Cloud. At Newfire, I optimized inference time of BERT-based models and built production NLP pipelines using spaCy and FlairNLP, with CI/CD managed through Jenkins. When projects require it, I handle the full pipeline including MLOps, backend development, and deployment, though my core strength lies in ML model development and optimization, alongside AI consulting. I also work directly with company leadership (CEOs/CTOs) to define AI implementation roadmaps and assess feasibility before costly investment.

Machine Learning Engineer (remote) October 2020 - September 2023
TIS Group

Key points:

- Developed ML systems in computer vision and NLP, including a healthcare solution (SENDD) for early detection of neurological deviations in infants.
- Led end-to-end development of Vineyard Angel - a drone imagery analysis system for vineyard health assessment and missing plant detection.
- Built ML pipelines and contributed to deployment efforts using Python (PyTorch, scikit-learn, OpenCV) on Microsoft Azure with Linux-based infrastructure.
- Participated in hackathons: built a scoliosis severity estimation model and CatAIog - an LLM-based PDF chatbot.
- Presented the Vineyard Angel project at the AI2FUTURE 2022 conference in collaboration with an agricultural partner.

Job description:
At TIS Group, I worked as a machine learning engineer across multiple projects, primarily in computer vision and natural language processing. One of the key projects was SENDD, a healthcare solution for early detection of neurological deviations in infants. I contributed across most of the ML pipeline, focusing on data science and model development, while supporting deployment efforts led by a teammate. The tech stack included Python (pandas, scikit-learn, TensorFlow, PyTorch), Jupyter, and Azure-based virtual machines running on Linux with Docker. We applied both classical (e.g., k-NN, SVM) and deep learning models (e.g., CNNs, autoencoders), and incorporated NVIDIA open-source technology in parts of the system. I also led the development of Vineyard Angel, a drone imagery analysis system that estimates missing vine plants and assesses canopy vigor. I handled everything from initial design to production deployment and later presented the project at the AI2FUTURE 2022 conference alongside our agricultural partner. In addition, I participated in hackathons where I led the development of a scoliosis severity estimation model and collaborated on building CatAIog, an LLM-powered chatbot that answers questions from PDF documents. I also created and voiced the video presentations for both solutions.

Software Engineer July 2019 - February 2020
Rimac Automobili

Key points:

- Developed infrastructure for autonomous driving systems, ensuring synchronization and health checks of vehicle sensors and camera inputs.
- Processed raw radar and sensor data in real-time using low-level C++ code.
- Collaborated with senior engineers to integrate systems into the vehicle's software stack.
- Utilized NVIDIA technologies and Git for development and version control.
- Balanced full-time and part-time responsibilities alongside university studies.

Job description:
At Rimac Automobili, I worked with the autonomous driving team, focusing on system infrastructure rather than direct machine learning development. My main responsibility was building and maintaining systems that ensured the synchronization and health of all onboard sensors, including cameras and radars. This involved real-time data processing and implementing logic to handle delays, missing frames, or timestamp mismatches. I worked extensively with low-level C++ code, integrating raw sensor data into the vehicle's software stack. My work also involved using NVIDIA technologies and maintaining code via Git. I collaborated closely with senior engineers throughout the process. This

role was a mix of full-time and part-time engagement, depending on my university schedule.

Software Engineer June 2016 - July 2016
COBE (Creators Of Beautiful Experiences)

Key points:

- Developed UI features in Unity using C# during a student engagement.
- Collaborated with a senior developer to align implementation with functional and design specifications.
- Contributed production-ready code and improved UI integration within an existing system.

Job description:
While still a university student, I worked as a freelance developer at COBE (Creators Of Beautiful Experiences). My main responsibility was building and integrating user interface components in Unity using C#. I contributed to an existing codebase, ensuring smooth integration of the UI with the system's functional core. Working under the supervision of a senior developer, I helped align the implementation with both design and functional specifications. This project was one of my first professional experiences in software development during my studies.

WEBSITE AND SOCIAL MEDIA

- Personal Website – Overview of my work, including selected projects and articles.
- LinkedIn Profile - Professional background, roles, and project involvement.
- GitHub Profile – Source code for personal and professional projects.

PUBLICATIONS

- AI Safety: A Quantitative Survey – Data-driven overview of subfields within AI safety.
- Whole-Body Vibrations Prediction – ML study on predicting vibrations in tractors; I led the ML component.
- Extraction of Human Preferences – RL-based AI Safety Camp project on implicit learning of human values. I contributed to RL environment setup, model tuning, and data analysis.

SIDE PROJECTS

- Transformer from Scratch – PyTorch implementation of the Transformer architecture from the ground up.
- SWEMentor – Multi-agent RAG AI system for automated code analysis and mentoring.
- Movie Script Generator – GPT-2-based system for generating extended, coherent screenplay fragments.
- Neural Network from Scratch – NumPy-based implementation of a basic neural network.
- Jobs for Juniors EU – Job board prototype built with Django, PostgreSQL, and AWS.

- MislavCoin – Personal ERC-20 token deployed to an Ethereum testnet; built using Solidity, Hardhat, and Truffle.
- RenaissanceinCroatia – YouTube channel with programming tutorials in Croatian, started in my teenage years.

PRESENTATIONS AND TALKS

- SWEMentor – A presentation of my SWEMentor side project.
- CatAIog – LLM-based PDF chatbot (TIS Group); I created and voiced the demo.
- Scoliosis Detection – Deep learning model for scoliosis severity estimation (TIS Group); demo created and voiced by me.
- AI2FUTURE 2022 Speaker – Presented Vineyard Angel, a CV system for drone-based vineyard analysis.

EDUCATION

M.Sc. in Computer Science 2018–2020
University of Zagreb, Faculty of Electrical Engineering and Computing

B.Sc. in Informatics 2015–2018
University of Zagreb, Faculty of Organization and Informatics

SKILLS AND TOOLS

- **Proficient in**
  - **Programming Languages:** Python
  - **ML libraries:** FastAPI, Hugging Face, NumPy, OpenCV, pandas, PyTorch, scikit-learn
  - **Tools:** Azure, Git, Jupyter, Linux
  - **Domains:** Computer Vision, GenAI, LLMs, NLP, Traditional ML
- **Familiar with**
  - **Programming Languages:** C++, SQL
  - **ML libraries:** Django, FlairNLP, Keras, PyInstaller, SciPy, spaCy, TensorFlow
  - **Tools:** AWS, CI/CD pipelines, Docker, GCP, Jenkins, Terraform, Vertex AI (incl. ADK), W&B
  - **Databases:** Chroma, CosmosDB, PostgreSQL
  - **Domains:** Agentic AI, ASR, Data Science, RAG, TTS

CERTIFICATES

- Generative AI with Large Language Models – DeepLearning.AI and AWS (20.02.2024.)
- DeepLearning.AI TensorFlow Developer - DeepLearning.AI (25.04.2020.)