

# OGNJEN JOVANOVIĆ

Software Engineer · Machine Learning · Computer Vision  
Bela Crkva, Serbia | (+381) 69 253 2088 | grmim.dev@gmail.com

---

## PROFILE

---

Newly graduated Software Engineering student with a strong foundation in computer science, mathematics, and applied machine learning. Experienced in building end-to-end ML systems—from dataset creation and CNN model development to GPU-accelerated training and real-time deployment on embedded platforms. Passionate about computer vision, autonomous systems, and edge AI. Seeking opportunities to contribute to teams working on challenging ML/AI engineering problems.

## EDUCATION

---

### **B.Sc. Software Engineering** 2022 – 2026

*Faculty of Economics and Engineering Management, Novi Sad, Serbia  
University Business Academy*

- Graduation thesis: "End-to-End Autonomous Vehicle Navigation Using CNNs" — designed and implemented a complete ML-driven autonomous navigation system with real-time embedded inference.

## PROJECTS

---

### **Autonomous Vehicle Navigation System** 2025 – 2026

*Graduation Thesis — End-to-End Deep Learning on Embedded Hardware*

- Built a complete autonomous driving pipeline: dataset collection (~20,000 labeled image-command pairs), PilotNet-inspired CNN training on NVIDIA GPU, dynamic quantization (FP32 → INT8, ~1.2x speedup), and TorchScript deployment on Raspberry Pi 4.
- Achieved real-time inference at ~20–25 Hz on ARM hardware, enabling autonomous hallway navigation using monocular vision with behavioral cloning.
- Integrated four HC-SR04 ultrasonic sensors as an independent safety layer with emergency obstacle avoidance, direction-aware recovery maneuvers, and a 500 ms watchdog timer.
- Developed a multi-threaded Python control server (`drive_server.py`) with dedicated threads for TCP control, telemetry, and MJPEG video streaming.
- Created a cross-platform Flutter mobile application with dual-joystick control, live camera feed, real-time sensor HUD, and seamless manual/autonomous mode switching.
- Tech stack: Python, PyTorch, TorchScript, OpenCV, Albumentations, Flutter/Dart, Raspberry Pi, GPIO, PWM.

### **GPU-Accelerated ML & LLM Research Infrastructure** 2025

*Personal Project — Local AI Research Environment*

- Architected a dedicated Ubuntu 24.04 workstation with NVIDIA RTX GPU acceleration for ML experimentation, LLM serving, and generative AI research.
- Deployed containerized ML infrastructure via Docker: Ollama, OpenWebUI, ComfyUI, and Lobe Chat with NVIDIA container toolkit passthrough.
- Configured CUDA/cuDNN toolchain and integrated PyTorch, TensorFlow, Keras, and OpenCV for end-to-end model development and benchmarking.
- Implemented a remote workstation-client workflow over LAN, eliminating cloud dependency for GPU-intensive research.

## Most Used Libraries for Machine Learning and AI<sup>2024</sup>

*Research Paper — Library Evaluation with Practical Demonstration*

- Authored a research paper surveying widely-used ML/AI libraries, analyzing their strengths, ecosystem fit, and industry adoption.
- Developed a companion Jupyter Notebook with an image recognition model trained on CIFAR-10 using PyTorch and NumPy with NVIDIA GPU acceleration.

## Trichinella Detection System (Upcoming)<sup>2026</sup>

*Independent Project — CNN Binary Classifier for Veterinary Diagnostics*

- Designing a Jupyter-based ML system for classifying trichinosis slide images as Trichinella-positive or negative, targeting  $\geq 95\%$  sensitivity for food safety applications.
- Architecture: EfficientNet-B3 with custom classifier head, Grad-CAM explainability overlays, Albumentations augmentation pipeline, and Docker-containerized GPU environment.

## Web Development Projects<sup>2024 – 2025</sup>

*Client Work — Full-Stack Websites*

- Designed and deployed production websites including [autizamjuznibanat.com](http://autizamjuznibanat.com) and [boggvart.com](http://boggvart.com) using Astro.js with contact form integrations (Formspree) and responsive design.

## TECHNICAL SKILLS

---

**Languages:** Python (advanced), C++ (advanced), Dart, Rust (basic), Java (basic), SQL

**ML / AI:** PyTorch, TorchScript, TensorFlow, Keras, OpenCV, Albumentations, NumPy, Pandas, Jupyter

**Techniques:** CNNs, Behavioral Cloning, Transfer Learning, Data Augmentation, Model Quantization, Grad-CAM

**Embedded / IoT:** Raspberry Pi, GPIO/PWM Control, Edge Inference, Sensor Integration, Real-Time Systems

**Infrastructure:** Docker, NVIDIA CUDA/cuDNN, Ubuntu Server, Git, SSH, Ollama, LLM Serving

**Web / Mobile:** Flutter, Astro.js, HTML/CSS/JS, Formspree

**Tools:** VS Code, Jupyter Lab, Linux CLI, systemd, Docker Compose

## EARLY ACCOMPLISHMENTS

---

### Petnica Science Center — Youth Talent Program<sup>2014</sup>

- Selected for a competitive scientific talent initiative among leading experts across disciplines. Completed a 14-day intensive program culminating in a research presentation on Graham's Scan algorithm and its implementation.

### Math, Physics & Programming Competitions

- Competed at municipal and regional level in mathematics, physics, and programming competitions during elementary school. Member of two science centers fostering early STEM interest.

## LANGUAGES

---

**Serbian:** Native

**English:** C1/C2 (Professional working proficiency)