

Ivan Mikhnenkov

AI Engineer & Researcher

Barcelona, Spain

ivanmikhnenkov.com

[LinkedIn](#) | [X](#)



6+ years in computer vision and generative AI. I take models from research and fine-tuning to efficient inference and production systems, with ownership of deployment and on-call operations.

EXPERIENCE

Ex-Human

Feb 2022 - Present

Lead Image Generation Engineer (Oct 2023 - Present)

Previously Computer Vision Engineer / Researcher

- Own the production AI system: **serving multiple models across distributed GPU resources, APIs and business logic.** Handle **on-call incident response**, monitoring and alerting.
- Fine-tune **image generation and editing models** for apps with **500K+ monthly active users and 30M+ images per month.** Build evaluation pipelines for image quality and identity consistency.
- Optimize **self-hosted diffusion and flow inference**, reducing latency and memory use while preserving quality. Avoid **tens of thousands of dollars per month in estimated API costs.** Accelerated video inference to **near real time** (MiniMax H3).
- Build **LLM-powered features** and use **vision-language models for moderation and quality estimation.**
- **Ship product features independently**, from prototype to deployment, and lead delivery across engineering and design. Use Amplitude and PostHog to guide iteration.

ZennoLab | Computer Vision Engineer / Researcher

Dec 2021 - Jul 2022

Built computer-vision services for browser automation and automated CAPTCHA solving.

RT Labs | Computer Vision Engineer

Feb 2020 - Jul 2021

Led satellite-based deforestation monitoring: designed custom neural architectures, trained and fine-tuned models, and built data and evaluation pipelines. Managed two engineers.

INDEPENDENT AI RESEARCH

- **LLM post-training:** implemented PPO, GRPO and DPO in PyTorch. GRPO raised Qwen3.5-0.8B's **GSM8K test accuracy from 12.8% to 57.1%** (1,319 questions; one run per method).
- **TinyDiT:** trained a **209M-parameter text-to-image transformer from scratch on 4.2M images** using flow matching and pretrained encoders. Built data pipelines, evaluation and interactive inference tools.

ENTREPRENEURSHIP

Built and tested AI product prototypes for video, multi-agent debates and e-commerce imagery.

BeVisionary (2018): built and launched my first product, a goal-setting and journaling app.

TECHNICAL SKILLS & EDUCATION

Tech stack: Python, PyTorch, FastAPI.

Infrastructure: AWS, Docker, Redis. **Monitoring & on-call:** Prometheus, Grafana, Sentry, PagerDuty.

Development tools: Claude Code, Codex, Cursor, Visual Studio Code.

Education: Bachelor's, Social Sciences (Sociology / Quantitative Methods).

Higher School of Economics, 2016-2020.

Languages: English (full professional proficiency), Russian (native).