

Alexey Zaytsev

Full-time learner. [Linux, Biology, Electronics] → AI.

<https://alexey.work/cv>
alexey.zaytsev@gmail.com
<https://github.com/xl0>

// FEATURED PROJECTS

lovely-tensors 1.4k ★ / **numpy** / **jax** Python, PyTorch, Num...
Tensor pretty-printing for human consumption. Widely adopted by AI researchers.

congusto-chat SvelteKit, TypeScript, PostgreSQL, Vercel AI S...
Open-source multi-provider chat interface, similar to OpenWebUI and t3chat.

lovely-docs SvelteKit, TypeScript, Python, LLM, MCP
AI-native documentation dehydrator for agentic access. CLI, web UI, and MCP server.

pelican SvelteKit, TypeScript, PostgreSQL, Vercel AI SDK
LLM-driven SVG/ASCII art generation gallery with iterative visual refinement.

latent-tools Python, PyTorch, ComfyUI, Diffusion Models
Custom ComfyUI nodes for fine-grained diffusion latent manipulation.

// EXPERIENCE

AI/ML Engineer (Contract)

July 2024 – Present

Various Clients

Some Client projects:

- AI Storyboarding Platform (Python backend): Integrated new image generation providers and refactored the generation abstraction layer across use cases (assets, scenes, inpainting).
- Earnings Call Analysis & CEO Training: Built a transcription pipeline (AWS Transcribe / Deepgram) with speaker diarization and LLM-driven financial metrics extraction. Companion CEO training module generates realistic analyst questions via TTS from call scripts, transcribes the CEO's spoken responses, and provides feedback on both substance and delivery (WPM, filler words).

Also built and maintain several open-source AI applications:

- Congusto Chat: Multi-provider multi-model AI chat UI offering ultimate flexibility. SvelteKit, Drizzle/PostgreSQL, Vercel AI SDK, Lucia auth, S3/CloudFront for media.
- lovely-docs: Documentation dehydrator for AI agents. Scrapes and clones codebases to generate recursive LLM doc summaries. SvelteKit UI, `npx` CLI, MCP server.
- Pelican: LLM-driven SVG/ASCII art gallery with iterative refinement (resvg + Jinja2 prompts). Client-

// RELEVANT SKILLS

AI/ML

Python, PyTorch, JAX, CUDA, Transformers, Diffusion Models, Computer Vision, LLM Agents, vLLM, Hugging Face, LangChain, RAG, Fine-tuning, LoRA, Stable Diffusion, CLIP, Embeddings, Vector Databases, Weights & Biases

Full-Stack

SvelteKit, TypeScript, FastAPI, PostgreSQL, Tailwind CSS, SQLAlchemy, drizzle-orm, Unified.js, Vercel AI SDK

DevOps

Docker, Linux, AWS, GitHub Actions, Supabase, Vercel, CloudFlare, S3

// EDUCATION

Université Denis Diderot

2013 – 2015

M.Sc. Systems and Synthetic Biology

LETI University

2003 – 2009

Specialist, Computer Security

// LANGUAGES

English	Professional
Russian	Native
Ukrainian	Elementary

side inference (Vercel AI SDK), SvelteKit, PostgreSQL, S3/CloudFront.

- latent-tools: PyTorch ComfyUI nodes for mathematical manipulation of diffusion latents (tensor blending, concatenation, noise generation) in SD/SVD.
- svelte-asciiart: Svelte 5 library rendering ASCII art as scalable SVGs. Computes precise viewBox constraints and bakes CSS styles for standalone exports.

SvelteKit · TypeScript · Vercel AI SDK · PostgreSQL · Drizzle ORM · Python · PyTorch · ComfyUI · MCP · LLM Agents · SVG · Tailwind CSS v4

Generative AI lead

Nov 2023 – July 2024

io.net

Led development of the generative AI platform supporting image, video, and 3D model generation.

Served tens of thousands of users, generating hundreds of thousands of images:

- Architecture & Inference: Built an end-to-end generation pipeline. The frontend talked to the backend that talked to Supabase/PostgreSQL, which queued tasks via RabbitMQ. Distributed GPU workers picked up the jobs, generated the media, uploaded to S3 (served via CloudFront+imgproxy), and pinged a backend webhook, which finally streamed updates to the client via Server-Sent Events.
- Model Fine-tuning: Extensive work with Stable Diffusion inference, custom model fine-tuning, and LoRAs to optimize generation quality and capabilities.
- Custom NSFW Detector: Off-the-shelf filters were unusable in production due to false-positive rates.

Trained a custom EfficientNet classifier on a large, intentionally generated dataset, achieving production-grade precision. Integrated directly into the GPU worker pipeline.

- MLOps & DevOps: Orchestrated and scaled the distributed inference infrastructure.

Stable Diffusion · PyTorch · RabbitMQ · PostgreSQL · Supabase · LoRA · EfficientNet · MLOps · SSE · S3

AI/ML Engineer (Contract)

Nov 2022 – Nov 2023

Various Clients

Independent contractor focused on AI apps and prototypes. Some that come to mind:

- Dog Finder: An experimental missing-dog app exploring both facial recognition and noseprint uniqueness. Trained EfficientNet with triplet loss.
- SAT Practice Generator: A dynamically updating test prep system seeded by real data. Built with Streamlit, Airtable, and OpenAI LLMs—migrating from GPT-3.5 to GPT-4 largely solved our initial correctness and consistency struggles.
- GPTEditor: An AI-assisted copywriting tool—like Cursor, but for prose. Users highlight a paragraph and instruct the LLM on how to rewrite it. Built with SvelteKit, Tiptap, and Supabase.

SvelteKit · OpenAI API · RAG · Supabase · Streamlit · TypeScript · Vector DBs · PyTorch · Computer Vision

Career Transition / Independent Study

May 2020 – Nov 2022

Self-directed

Deep learning and AI intensive study period focused on understanding the fundamentals:

- lovely-tensors: Built an open-source PyTorch debugging library for human consumption, which saw wide adoption among AI researchers. Later expanded to lovely-numpy and lovely-jax.
- GPT-2 from Scratch: Built and trained a GPT-2 style transformer using raw PyTorch.
- Autograd on NumPy: Implemented a custom autograd library on top of NumPy, featuring enough operators to train a miniature GPT-2 variant on the TinyStories dataset.

PyTorch · Deep Learning · NumPy · JAX · Transformers · Autograd · Python · Open Source


Electronics Engineer

Dec 2018 – May 2020

Okra Solar — Canberra / Phnom Penh / Shenzhen

Smart solar grid tech: designed the main grid hub (Solar → Battery → regulated DC loads and high voltage sharing), a debugging adapter, and auxiliary internal boards. Schematics, PCB design, DC-DC converters and analog monitoring circuits. Led a small hardware team (3 engineers).

Analog Electronics · Power Electronics · Maximum power point tracking · KiCad




Electronics and Prototyping Engineer

Sept 2016 – Dec 2018

Self-employed — Shenzhen, China

Dozens of client projects across digital, analog, and power electronics — from initial concept to production-ready designs and manufacturing oversight. Clients in medical devices, scientific instruments, food and consumer electronics. Designed and set up end-of-line automated testing, trained factory workers, and oversaw manufacturing in Shenzhen factories.

Digital Electronics · NFC · E-Ink · Bluetooth Low Energy · Sub-1G radio · Embedded systems · Power electronics · Motor controller · Synchronous rectification · DC-DC converters · Buck · Boost · LED display · Analog Electronics · Differential amplifiers · EKG · Electrostimulation · EEG · Electrowetting · Digital microfluidics · Autodesk Inventor · Fusion 360 · KiCad · Altium Designer



CTO & Co-founder

Sept 2015 – Sept 2016

Koi.science — Paris / Shenzhen

EdTech startup providing affordable laboratory equipment for molecular biology education in schools. Designed spectrophotometer, electrophoresis and transilluminator — hardware, electronics, mechanical enclosures, and embedded software. Co-founded with fellow students from the CRI Paris program.

Electronics Engineering · Mechanical Engineering · Microcontrollers · Embedded programming · Autodesk Inventor · Fusion 360 · KiCad




Research Intern

May 2015 – Sept 2015

Centre de Recherches Interdisciplinaires (CRI Paris) — Paris

Developed the first iteration of a low-cost, open-source spectrophotometer — the project that later evolved into the Koi.science startup.

Electronics Engineering · Mechanical engineering · Embedded programming · Spectrophotometry · KiCad · Education outreach · Autodesk Inventor



Research Intern


May 2014 – Jan 2015

Institut Pasteur — Paris

Characterizing bacteriophage origin of replication region, exploring the potential of applying genetically engineered λ -phages to antibiotic-resistant bacterial infections in burn victims.

Contributed to a published paper.

Bacteriophages · λ -phage · Staphylococcus Aureus · PCR · Plasmids · Molecular Biology · Electrophoresis · CRISPR-Cas9 · Bioinformatics · BioPython · High-throughput sequencing · Gibson Assembly



Research Intern

Jan 2014 – April 2014

Institut Curie — Paris

Researching the influence of electric fields on mammalian cell mobility and wound healing.

Microfluidics · Protein micropatterning · PDMS · Electrophoresis · Light Microscopy · Mammalian cell culture · Electronics prototyping



Sr. Software Engineer

June 2011 – July 2012

Nexenta Systems

Virtio drivers for network and block devices, general kernel maintenance, bugfixes and development on illumos (OpenSolaris fork).

Device drivers · Virtio · ZFS · Networking · Block storage · Solaris · C

Linux Kernel Engineer-Programmer

Jan 2005 – Jan 2010

AltEIl and Protei ltd Saint-Petersburg, Russia

Linux kernel development for high-bandwidth telecom equipment. AltEIl: 10 Gbps network routers. Protei: telephony over dedicated lines, WiMAX and GPRS/UMTS base stations. Device drivers, bootloaders, and board support packages on MIPS64, ARM and x86.

Device drivers · Bootloaders · Networking · ASIC · Board Support Packages · Linux Kernel · C · Assembly · ARM · MIPS · Intel XScale

