Saulius Beconis

Applied ML Engineer · evaluation, agents, audio

Vilnius, Lithuania (UTC+2) · remote · saulius@sbeconis.com · saulius.beconis@gmail.com · +370 696 65675
github.com/sauliusbeconis · linkedin.com/in/sauliusbeconis · sbeconis.com/lab

Applied ML engineer with a research habit: held-out splits, negative controls, and results written up even when they come out negative. My thesis work compares learned and classical representations of rat pup vocalizations across held-out litters. I also run AxonGate, a paid web-access API for AI agents with an MCP server, and publish measured experiments on document AI, retrieval and local LLMs at sbeconis.com/lab.

RESEARCH AND PROJECTS

Thesis research, Vilnius University · deep learning on rat pup ultrasonic vocalizations · 2019 – present

Python · PyTorch · scikit-learn · librosa · mixed models · uv · pytest · mypy

- Designed the recording protocol and a sound-isolated ultrasonic rig; recorded 588 sessions (49 h) from 28 tracked pups over a 21-day developmental window (ethics permit G2-20).
- Call detector hand-checked at ≈ 0.93 precision (≈ 0.99 above 14 dB peak SNR) and ≈ 0.99 recall.
- Compared a VAE, a spectrogram CNN and a self-supervised model with classical acoustic features, leave-one-litter-out: per-pup age error fell from 3.42 to 2.38 days with the CNN.
- Rejected the self-supervised model because its background-only embeddings also predicted age. Negative controls also caught a session leak in leave-one-animal-out splits.
- Contour shape explains ~76% of the VAE gain within the cohort; call order adds nothing beyond single calls. Both results are written up.
- Pre-registered a discovery and confirmation round across held-out litters; 21 decision records, every run logs commit, config, split and seed.

S. BECONIS Lab · 26 published applied-AI experiments

sbeconis.com/lab

- Invoice extraction with local Qwen3-VL-8B: 95.5% straight-through on 400 synthetic invoices in four languages. Seven plain-code checks took silent errors on a fresh set from 4.8% to 0%, at 1.3% false rejects.
- Local coding agent on 40 QuixBugs programs: with tests available it still claimed 8 false fixes. Letting only passing tests end a task cut false fixes to 0; Qwen3-14B then fixed 82.5%.
- Compared keyword, vector, hybrid and reranked retrieval on public test sets; measured local LLMs on one 12 GB GPU.

AxonGate · pay-per-call web API for AI agents · 2026, live

Python · FastAPI · Redis · MCP · x402 · Railway
api.axongate.one · github.com/sauliusbeconis/AxonGate

- Page reading, search, grounded answers and claim verification for agents. Quotes are checked verbatim against the source, and verdicts without proof are downgraded.
- MCP server, OpenAPI, llms.txt and agent card; 31 paid calls from 11 outside wallets in about 200 days live.
- Every fetched page is scanned for prompt injection; private and metadata addresses are refused.

AxonTrade · futures research and execution lab · 2025 – 2026

github.com/sauliusbeconis/AxonTrade

- Walk-forward tests, frozen holdouts, Monte Carlo shuffles and slippage stress on ~800,000-row order-flow data; dropped adaptive parameter selection when it failed out-of-sample.

iGEM Vilnius-Lithuania · IT team lead · 2019

- HMM profile search over raw MG-RAST metagenomes against a user-selected genome profile. The team won a gold medal.

EXPERIENCE

Generative AI pipelines · freelance · 2023 – 2026

Image and video generation workflows, tooling and training for a private client (ComfyUI, local models, Python).

Videographer and editor · staff and freelance · 2021 – 2025

Shoots, editing and graphics; staff videographer at Bored Panda (2022 – 2023), video consultant to the National Library of Lithuania (2021 – 2022).

QA Engineer · NSoft · 2020 – 2021

Software and hardware testing, API testing in Postman, regression runs.

EDUCATION

Biology, Vilnius University · Life Sciences Center · 2016 – 2020

Bachelor's thesis in progress (the vocalization research above), defence by June 2027.

SKILLS

ML PyTorch, scikit-learn, NumPy, pandas, librosa, VAEs, CNNs, mixed models

Evaluation group-held-out splits, negative controls, pre-registration, error analysis

LLM llama.cpp, local vision-language models, RAG, MCP, prompt-injection screening

Engineering Python, FastAPI, Redis, Linux, Docker, uv, pytest, mypy, GitHub Actions

Spoken Lithuanian (native), English (C2)

OTHER

Chairman, Society of Innovative Medicine (2018 – present), a registered association linking businesses and scientists in medicine.