Uri Maayan

Computational Researcher & Algorithm Developer

Givatayim, Israel • 054-3032928 • uriuriuri7@hotmail.com • LinkedIn • GitHub

M.Sc. Chemical Engineer (93.75 / 100) turned independent AI engineer. Combines a strong numerical-algorithms background — including a first-of-its-kind method for damping-coefficient extraction in nanotribology — with hands-on experience building production-grade AI systems end to end. Over the past year designed and shipped five independent Python systems spanning real-time RAG, computer-vision detection, multi-LLM orchestration, live broadcast automation, and ML-systems work porting a state-space TTS model to AMD / ROCm. Bilingual (Hebrew / English), fast learner, autodidact.

SELECTED INDEPENDENT PROJECTS

EventQuote AI — Real-Time Event-Production Quotation System

- Solo-built SaaS that listens to live client meetings (Hebrew or English), extracts technical requirements with an LLM, matches them to an equipment inventory via RAG, and generates a ready-to-send .docx / .pdf quote in seconds.
- Python · FastAPI · Vue 3 + TypeScript + Tailwind · ChromaDB · SQLAlchemy + Alembic + PostgreSQL · Whisper live chunked transcription · OpenAI / Gemini / Ollama · Docker. 16 REST modules, 1360+ pytest tests, 235-key Hebrew / English i18n with full RTL, multi-tenant JWT auth.
- WebSocket live-transcription endpoint, analytics dashboards (ECharts), client approval portal, WhatsApp Business bot, QuickBooks / Hashavshevet / Priority ERP accounting integrations.

SkyWatch — Multi-Camera Missile-Detection System

- Real-time visual-awareness system aggregating public camera feeds across Israel to automatically detect missile launches and Iron-Dome intercepts during rocket attacks.
- Python · ffmpeg with hardware decode (D3D11VA / VideoToolbox) · yt-dlp · OpenCV · UniFi Protect local API · Pikud HaOref alert integration.
- 5-stage CV pipeline: ROI crop → adaptive background subtraction → streak / flash signature detection → multi-frame tracking → false-positive filtering. Cross-camera temporal correlation, 3×3 grid display with auto-focus on the detecting camera, rolling pre-buffer recording, separate day / night algorithms.

Ozer-AI — Multi-LLM Task-Orchestration Middleware

- AI middleware that delegates coding tasks to free LLM providers using a decentralized Agent Teams architecture — saving senior-model tokens for planning and review while workers handle implementation.
- Python 3.12 · asyncio · httpx · Pydantic · Click + Rich CLI. Providers: DeepSeek, OpenRouter, Groq, Alibaba Qwen, Gemini, Puter.js.
- DAG-based plan execution with dynamic task claiming, inter-task context sharing so dependent workers see upstream output, role-matched routing (backend / test-writer / frontend / architect), LLM cross-review with automatic retry and escalation, and a final LLM-powered integration-synthesis pass.

Channel Shesh — AI Hebrew Co-Host Livestream System

- End-to-end autonomous livestream system that monitors rocket alerts and social-media content in real time, generates Hebrew commentary with Gemini, synthesizes speech, and drives OBS automatically — effectively an AI anchor.
- Python · FastAPI producer dashboard · OBS WebSocket v5 · Gemini (content analysis + script generation) · Edge TTS · ChromaDB (long-term episodic memory) · Redis (short-term memory). 114 pytest tests.
- Pipeline: Pikud HaOref poller → content batcher → Gemini fact extractor → Hebrew scriptwriter → TTS synthesis → OBS scene / source controller → pacing engine, orchestrated by a state machine with pluggable content sources.

Zonos-Hebrew on AMD / ROCm — Pure-PyTorch Mamba2 + Attention Backbone

- Ported the Zonos state-space Hebrew TTS model to AMD GPUs (RX 6800 XT) by rewriting the hybrid Mamba2 + Attention backbone in pure PyTorch — removing all CUDA-only dependencies (mamba_ssm, causal_conv1d, flash_attn). Weight layouts match upstream exactly, so the pretrained checkpoint loads unmodified.
- Implemented the SSD chunked scan, RMSNormGated, and single-step decode path from scratch. Identified a bf16 SSM state-drift bug that collapsed audio past ~2.8 s of generation and forced the recurrence to fp32; matched flash_attn's half-split rotary-embedding convention (not interleaved), which is what the checkpoint was trained against.
- ROCm workarounds: diagnosed a memory-access fault in ROCm SDPA's flash / efficient-attention backends on long KV sequences and forced the MATH backend; manually expanded KV heads to sidestep an enable_gqa=True ROCm bug; replaced torch.roll with in-place shifts in the conv cache to remove per-step allocator churn.

RESEARCH & PROFESSIONAL EXPERIENCE

Graduate Researcher (M.Sc. Thesis) | Berkovits Lab, Ben-Gurion University | Be'er Sheva | 2021 – 2024

- Thesis: "Extracting damping coefficients from nanoscale friction slip dynamics using linear approximations." Developed the first direct algorithmic method for extracting damping coefficients from AFM slip dynamics — no prior technique existed for quantifying this parameter from molecular-dynamics trajectories. Thesis grade: 94 / 100.
- Built a Python SDE simulation engine from scratch: Langevin-equation integrator using 4th-order Runge-Kutta with Gaussian noise, Numba-JIT acceleration, 1-D and 2-D systems. Generated and processed 16 M+ trajectory data points across parameter sweeps.
- Designed a novel slip-detection algorithm based on windowed moving standard deviation (exploiting rapid force-value changes during slip), followed by automatic segmentation and nonlinear fitting of a spring-damper equation of motion to each isolated slip. Repository: github.com/Zuzutus/sde-solver.

Teaching Assistant & Tutor | Ben-Gurion University | Be'er Sheva | 2021 – 2023

- Chemical Engineering Laboratories: instructed undergraduate lab sessions, explained the underlying physics and chemistry, graded detailed lab reports, verified pre-lab readiness via quizzes, and graded homework for the Control Systems course.
- Dean's Office Students-With-Disabilities Program: tutored 2 – 3 undergraduates individually in Calculus II and General Chemistry, adapting explanations to each student's learning profile.

Non-Commissioned Officer | IDF (Israel Defense Forces) | Israel | 2.5 years

- Managed a unit of security soldiers guarding a central Israel base: patrol coordination, shift handovers, and personnel management. Completed the NCO course.

EDUCATION

M.Sc. Chemical Engineering (Thesis track) | Ben-Gurion University of the Negev | 2021 – 2024

- Weighted final grade: 93.75 / 100. Thesis grade: 94 / 100.
- Relevant coursework: Deep Learning for Physical Systems & Inverse Problems / PINNs (91, taught in English), Advanced Thermodynamics (93), Analytical Mathematical Methods in Chem. Eng. (92), Nanomaterials (95), Ceramic Membranes (100), Alternative Energy Sources (96).

B.Sc. Chemical Engineering (Accelerated Track) | Ben-Gurion University of the Negev | 2018 – 2021

- Completed a four-year program in three years after convincing the faculty committee to approve an accelerated track.

Pre-Academic Studies — Mathematics & Programming | The Open University of Israel | 2016 – 2018

- Dean's List. Strong performance in Calculus I / II, Linear Algebra (for mathematicians), and introductory Java programming.

TECHNICAL SKILLS

- **Languages:** Python (expert), MATLAB, TypeScript / JavaScript, SQL, Scala, Java, C / C++.
- **AI / ML:** LLM orchestration (OpenAI, Gemini, DeepSeek, Groq, Ollama), RAG pipelines, ChromaDB, PyTorch (incl. custom Mamba2 / SSD / rotary-embedding / SDPA backends), TensorFlow, Physics-Informed Neural Networks, Whisper, Hebrew TTS (Zonos, Edge TTS), local-LLM deployment on AMD ROCm / Vulkan.
- **Backend & Data:** FastAPI, SQLAlchemy + Alembic, PostgreSQL, asyncio, WebSockets, pytest, NumPy, SciPy, pandas, Numba, stochastic differential equations, Monte Carlo, Runge-Kutta, time-series analysis, large-scale data processing.
- **Frontend & Media:** Vue 3 + Vite + TypeScript + Tailwind CSS, Pinia, vue-i18n (Hebrew RTL + English), ffmpeg hardware-accelerated decode, OpenCV, OBS WebSocket automation, RTSP / HLS / MJPEG ingest, yt-dlp.
- **Tooling & Analytics:** Docker + docker-compose, uv, bun, Git, Linux / WSL2, GA4, Looker Studio.

PUBLICATIONS & CONFERENCES

- E. Chetrit et al., "Nonexponential kinetics captured in sequential unfolding of polyproteins over a range of loads," Current Research in Structural Biology 4 (2022), pp. 106 – 117.
- Poster presentation, NANO.IL — annual national nanotechnology conference.

LANGUAGES

Hebrew (native) · English (fluent, self-taught from childhood)