Mariusz Kurman, MD

LLM Research Engineer

Pretraining · Model architectures · Post-training · Clinical AI

mariusz.kurman@gmail.com

github.com/mkurman · huggingface.co/mkurman · LinkedIn

Physician and ML engineer building language models from architecture and training code to released checkpoints and inference integrations. Pretrained custom models on 144B and 250B tokens using TPUv4; implemented JAX/Flax and PyTorch models, training pipelines, and teacher-feedback GRPO. Eight years building and leading applied AI products and teams, with a clinical background grounded in authoring and reviewing patient simulations.

SELECTED MODEL TRAINING & RESEARCH

NeuroBLAST v3

Hybrid architecture & TPU pretraining | Code · Checkpoint

- Designed and implemented a 597M-parameter, 72-layer language model in JAX/Flax and PyTorch, combining attention, dilated causal 2D convolution, and long-range residual connections.
- Pretrained on approximately 144B tokens of PleIAs/SYNTH using TPUv4-16 through Google TRC; released BF16 weights and custom Hugging Face modeling code.
- Built training code with gradient accumulation, streamed data loading, Optax optimization, Orbax checkpoint save/restore, and Weights & Biases logging.

ConvGPT

Compressed decoder & inference integrations | Code / serving · 250B-token checkpoint

- Designed a convolutional compression architecture mapping 1,296-dimensional embeddings to a 144-dimensional decoder stream; pretrained on approximately 250B tokens using TPUv4-16.
- Implemented custom vLLM and SGLang model modules, including attention/KV-cache integration and checkpoint weight loading.
- Extended the research in ConvGPT-v2 with causal 1D/2D convolution, sparse chunk-token retrieval, custom Triton kernels, and configurable training/evaluation pipelines.
- Executed an eight-variant, time-budgeted architecture screen with separate training/evaluation files, and five 5,000-step tokenizer comparisons.

Teacher-feedback post-training

GRPO & supervised fine-tuning | Training code · Released model

- Built a student/teacher training workflow with sampled completions, teacher-scored rewards, group-normalized advantages, reference-policy KL regularization, and optional SFT on feedback.
- Released a Llama-3.2-3B-derived reasoning checkpoint after 1,500 teacher-feedback GRPO steps; published model usage instructions and training code.

ClinicalGym & ClinicalBench

Clinical agent environments & evaluation | Environment · Benchmark

- Built an API-first clinical simulation and evaluation layer over 1,300+ physician-authored cases, each validated by at least two other physicians; personally authored approximately 300 cases.
- Implemented typed clinical actions, replayable trajectories, diagnosis-gated scoring, and filtered preference/reward-model dataset exports for downstream training.

PROFESSIONAL EXPERIENCE

Kruk SA

AI/ML Tech Lead, R&D · Sep 2023 – Present

- Lead a four-person R&D team developing and deploying LLM document-understanding systems that process thousands of documents each month.
- Own delivery from technical evaluation and prototyping through production deployment and MLOps monitoring across cloud environments.

MedIT Solutions

Founder, CEO & AI/ML Tech Lead · Apr 2018 – Oct 2026

- Built and led a 36-person cross-functional engineering, medical, and content team while remaining hands-on in model development.
- Shipped Medcases, an interactive clinical-case learning platform serving 55k+ users; built medical knowledge-extraction and generation workflows for educational content.
- Trained and released open models including MedIT One, a 140M-parameter LM pretrained from scratch with a published 9B-token checkpoint; also released Llama-3.2 SUN chat models and Polish medical model merges.

Medcases · MedIT One checkpoint

Research.com

Machine Learning Engineer & Project Owner · Mar 2021 – Jun 2023

- Built ML/NLP systems for scholar and venue ranking and author disambiguation across millions of papers and researchers.
- Owned the technical roadmap and coordinated cross-functional delivery from research and model development to platform release.

Earlier software engineering

Java & PHP Developer · 2007 – 2018

- Developed software across multiple companies and products, including edoktor24.pl and x-kom.pl.

ADDITIONAL OPEN-SOURCE ENGINEERING

SynthLabs — Built open-source tooling for generating SFT, reasoning, and tool-call datasets, including clinical training data.

Zorai — Built a Rust daemon for persistent agent execution, memory, tools, approvals, and long-running goals across GUI, TUI, CLI, and MCP clients; 324 GitHub stars.

TECHNICAL EXPERTISE

Model training: Pretraining from scratch, SFT, teacher-feedback GRPO, synthetic data, preference-data construction; BF16, gradient accumulation/checkpointing.

Frameworks & systems: Python, PyTorch, JAX/Flax, Hugging Face Transformers/TRL, Optax, Orbax, Triton; TPUv4, vLLM, SGLang.

Production & evaluation: Weights & Biases, MLflow, Kubeflow, OpenShift AI, FastAPI; AWS, GCP, Azure; clinical trajectory analysis and evaluation.

EDUCATION & LANGUAGES

MD, Medicine · University of Warmia and Mazury in Olsztyn · 2011 – 2017

BEng, Computer Science and Electronics in Medicine · Warsaw University of Technology · 2007 – 2011

Polish (native) · English (C1)