

Sukhrob Ilyosbekov

Machine Learning Engineer · Computer Vision & LLM Systems

📍 Portland, ME 📞 (857) 867-1942 ✉ silyosbekov@gmail.com

🌐 suxrobgm.net/research 🎓 [Google Scholar](https://scholar.google.com/citations?user=suxrobgm) 🔗 linkedin.com/in/suxrobgm 🐙 github.com/suxrobgm

SUMMARY

Machine learning engineer who takes models from research question to production service. Trained in computer vision and multimodal learning, with published work in medical imaging and representation learning, on top of nine years of building software that people depend on. Works across the whole loop: data, training and evaluation, serving, and the monitoring that follows a launch. Applied-math background.

TECHNICAL SKILLS

Deep Learning & CV: PyTorch, TensorFlow, CNNs, Vision Transformers, DINOv3, EfficientNet, YOLO, contrastive learning, GradCAM++, OpenCV

LLM Systems: Claude/GPT/DeepSeek APIs, tool-use agents, MCP, RAG, vision-language models, Hugging Face

Math & Modeling: Probability and statistics, linear algebra, optimization, time-series forecasting, feature engineering, A/B testing, error analysis

Languages: Python, C#, C++, TypeScript, Kotlin

Data & Serving: NumPy, pandas, scikit-learn, FastAPI, PostgreSQL, SQL Server, MongoDB, Redis, OHIF/DICOM

Infrastructure: CUDA, mixed-precision and multi-GPU training, model serving and drift monitoring, Docker, Kubernetes, GitHub Actions, AWS, Azure

Software: ASP.NET Core, SignalR, Blazor, Node.js, Bun, NestJS, React, Next.js, Angular

EDUCATION

Northeastern University, M.S. in Computer Science, GPA 4.0, Boston, MA

Jan 2024 – May 2026

Suffolk University, B.S. in Computer Science, Cum Laude, Boston, MA

Sep 2021 – May 2023

Honors: National Physics Olympiad, 1st Place (Uzbekistan) · Associate Member, Sigma Xi · University Coding Competition Winner, TUIT

PROFESSIONAL EXPERIENCE

Senior AI/ML Engineer | [EmTech Care Labs](#) | Portland, ME

Jan 2025 – Present

- Decline-risk model over daily activity and care-log data for Alzheimer's residents; 0.81 AUC on held-out data.
- NLP pipeline extracting symptoms, medications, and care events from free-text clinical notes; 0.90 F1.
- ML serving on de-identified data under HIPAA: nightly retraining, drift monitoring, encrypted PHI, audit logging.
- Lead engineer on the platform itself (TypeScript monorepo, AWS); passed an outside HIPAA readiness review with no critical findings.

Machine Learning Engineer | [AllFactors](#) | Santa Clara, CA

Oct 2021 – Dec 2024

- Property-valuation and market-forecast models over 2M listings in 40+ metros; median error 9% to 5% against the rules-based baseline.
- Real-time feature pipeline on SignalR and SQL Server change tracking; nightly retraining on the day's transactions.
- Served predictions into Blazor WebAssembly dashboards, with automated tests on the data and prediction flows.

Earlier: Software Engineer, Smart Meal Service, Russia (2020 – 2021), C#/.NET control software for food kiosks and a robotic cashier. Game Developer, Pentallight Technology, Malaysia (2020 – 2021), Unity networking and hand tracking for a VR simulation. Freelance Developer (2019 – 2020), Python NLP and .NET. Game Developer, EC Dev Team, Uzbekistan (2016 – 2019).

SELECTED AI/ML PROJECTS

MorphoCLIP

[GitHub](#) | [Paper](#)

PyTorch, DINOv3, BioClinical ModernBERT, contrastive learning, CPJUMP1

- Contrastive image-text model matching Cell Painting images of drug- and gene-perturbed cells to plain-language treatment descriptions.
- Frozen DINOv3 and BioClinical ModernBERT backbones with trained projection heads; trains on one consumer GPU.
- Evaluated on CPJUMP1 (51 plates, 3M+ images). First author; led design, training, and evaluation. On arXiv.

Localize, Don't Beautify

<https://arxiv.org/abs/2608.02841>

Image-editing APIs, inpainting, ArcFace, prompt engineering

- Benchmarked prompt steering, masked compositing, and inpainting across six commercial image editors on 196 cosmetic-surgery preview edits.
- Scored identity preservation (ArcFace) and edit localization; masked compositing localized best. Sole author. On arXiv.

MelanomaNet

[GitHub](#) | [Paper](#)

PyTorch, EfficientNet V2, focal loss, GradCAM++, OpenCV, ISIC 2019

- EfficientNet V2 with focal loss over nine ISIC 2019 categories, 25K images; 0.86 weighted F1.
- GradCAM++ attention scored against ABCDE dermatology criteria measured from the lesion mask. Sole author. On arXiv.

Med Image Scanner

<https://github.com/suxrobgm/med-image-scanner>

FastAPI, PyTorch, OpenCV, OHIF Viewer, DICOM, Next.js

- Pulls X-ray, CT, and MRI from hospital imaging systems over DICOM and runs PyTorch detectors: pneumonia on chest X-ray, hemorrhage on head CT.
- Renders predictions as overlays inside the OHIF viewer; on-the-fly de-identification, audit logging, and role-based access.

LogisticsX

<https://logisticsx.app> | [GitHub](#)

Claude API, MCP, .NET 10, EF Core, SignalR, Angular 21, Kotlin Multiplatform, PostgreSQL

- Claude tool-use agent that matches freight loads to trucks, checks federal hours-of-service, and plans multi-stop routes without a dispatcher.
- Full product around it: multi-tenant .NET backend, Angular portals, Kotlin Multiplatform driver app, load-board and telematics integrations.

Bookshelf Scanner

<https://github.com/suxrobgm/bookshelf-scanner>

YOLO, Moondream2 VLM, llama.cpp, FastAPI, Angular

- YOLO segmentation isolates each book spine; Moondream2 reads title and author; Angular UI for corrections and export. Course Outstanding Project Award.

OGStack

<https://ogstack.dev> | [GitHub](#)

Next.js 16, ElysiaJS, Bun, Satori, Cloudflare R2, Stripe

- Multi-model layer (Claude, GPT, DeepSeek) that reads a page's branding and generates its social-share image; 10K+ images in beta, served from a global CDN.