Rustam (Russ) Mukhtarov

San Francisco, CA | (628) 358-6989 | rmukhtarov@sfsu.edu | Website |GitHub | LinkedIn

PROFESSIONAL SUMMARY

ML/AI Engineer and M.S. candidate at SFSU with 3+ years building and deploying production ML pipelines. Proven results: 98% model accuracy and 20% faster deployment at RevSolz Corp; active thesis on 3D deep learning for medical imaging (PyTorch, MONAI, U-Net). Seeking mid-level ML/AI Engineer or Data Scientist roles in the SF Bay Area or remote.

SKILLS

ML & AI: PyTorch, TensorFlow, Scikit-learn, MONAI, YOLO, OpenCV, Transfer Learning, CNNs, 3D U-Net, Deep Learning
Data Science: Python, pandas, NumPy, SciPy, Matplotlib, SQL, Statistical Modeling, Data Preprocessing, A/B Testing
LLM & GenAI: LangChain, LlamaIndex, Hugging Face, OpenAI API, Gemma, RAG Pipelines, FAISS, Chroma, Prompt Engineering, Fine-tuning (LoRA/QLoRA), Conversational AI, REST API integration
MLOps & Cloud: Docker, Git, Bash, Linux, model deployment, pipeline automation — AWS/GCP
Robotics & CV: ROS, Gazebo, DroneKit, VTK, VMTK, Raspberry Pi, MATLAB, Ansys
Full-Stack: JavaScript, TypeScript, Python, C++, Flask, React, Next.js, Node.js, Tailwind, PostgreSQL, HTML/CSS, PyQt5
Languages: English (Fluent), Russian (Fluent), Turkish (Fluent), Azerbaijani (Fluent), Spanish (Beginner)

WORK EXPERIENCE

AI-Assisted Medial-Axis Modeling of the Human Fallopian Tube — San Francisco, CA
Advisor: Prof. Kazunori Okada, SFSU — Aug 2025 – Present

- Developing a **3D deep learning pipeline** (PyTorch, MONAI) for automated segmentation of micro-CT scan data; targeting ≥0.85 Dice score in lumen segmentation using a custom 3D U-Net architecture.
- Applying medial-axis extraction algorithms to generate centerlines with curvature, radius, and torsion descriptors using **VTK, VMTK, and SciPy**.
- Building a Python-based quantitative shape analysis toolkit for segment classification and biomechanical modeling; dataset designed to be simulation-ready and reproducible.

RevSolz Corp. — Calgary, AB, Canada (Remote)
AI Team Lead — *Jul 2020 – Jul 2023*

- Led a team of 4 data science interns to build and deploy ML pipelines for oil industry operational data, achieving **98% accuracy** in efficiency identification and improving company-wide decision-making.
- Designed end-to-end ML pipelines covering data preprocessing, feature engineering, model training, and deployment — reducing deployment time by **20%** and improving production readiness.
- Built and maintained supervised learning models (classification, regression) using **Python, Scikit-learn, and TensorFlow**; evaluated performance with cross-validation, precision/recall, and F1 metrics.
- Mentored interns in ML techniques and best practices; 2 interns were offered part-time roles based on their contributions.
- Communicated results, model insights, and technical trade-offs to non-technical project managers through clear documentation and presentations.

EDUCATION

San Francisco State University — San Francisco, CA
M.S. in Artificial Intelligence & Data Science — *Expected Dec 2026*
Teaching Assistant — Theory of Computing — *Aug 2024 – Dec 2024*

- Supported 30+ students in automata, formal languages, and computational complexity; assisted with grading and course materials.

Yildiz Technical University — Istanbul, Turkey
B.S. in Mechatronics Engineering — *Jun 2023*

PROJECTS

RAG Document Assistant - GitHub *July 2026*

- Built a production-style retrieval-augmented generation system for domain PDFs: FastAPI backend, local vector search (ChromaDB), and schema-constrained LLM generation with per-claim citations resolved server-side to source file and page.
- Designed a three-layer hallucination guard (grounding prompt contract, JSON-schema output enforcement, citation-index validation); benchmarked on 20 questions over a 70-page technical report: 0.87 retrieval hit rate, 100% abstention on unanswerable questions, zero fabricated facts.
- Built an evaluation harness and retrieval-inspection tooling; diagnosed reproducible retrieval failures to root cause (chunk dilution, embedding relevance bias) and designed the mitigation (hybrid BM25 + dense retrieval).
- Shipped with 29 offline tests using injected fake LLM/embedding clients, CI on every push, and one-command Docker Compose deployment.

Medical Image Disease Classifier API- GitHub *June 2026*

- Built and deployed a production-style medical image classification REST API (FastAPI + PyTorch) that returns a predicted class, full probability distribution, and Grad-CAM explanations from an uploaded scan, with input validation and structured error handling.
- Engineered a reproducible training/inference pipeline with shared preprocessing and a self-describing model checkpoint to eliminate train/serve skew; reached 94.1% accuracy / 0.94 macro-F1 on a reproducible synthetic 3-class benchmark (non-clinical).
- Containerized the service with Docker and configured GitHub Actions CI running 20 automated tests (API, preprocessing, model, inference) plus linting. Designed a plug-in ImageFolder interface so real datasets swap in with no code changes; documented model limitations and non-clinical scope.

AI Chatbot Web App *Jan 2025 – Present*

- Built a full-stack conversational AI application integrating **Google's Gemma LLM** via the Generative AI API; implemented prompt engineering for tone control and dynamic response generation.
- Developed Flask backend with REST API architecture; custom HTML/CSS/JS frontend optimized for conversational flow, emotional tone feedback, and smooth human-AI interaction.
- Iterated on UX based on user testing — refined conversation pacing, response latency, and visual cues.

DriveShare App & Platform - GitHub *Jan – May 2025*

- Co-built a peer-to-peer driveway rental platform using **Next.js, Tailwind, and PostgreSQL**; developed full product flow including listings, location-based search, bookings, and payments.
- Designed UI/UX in Figma; conducted iterative usability testing and refined driver and homeowner dashboard flows based on real user feedback.

Alzheimer's Disease Classification - GitHub *Aug – Dec 2024*

- Fine-tuned **VGG16, InceptionV3, and ResNet50** transfer learning models on PET scan data for multi-class Alzheimer's stage classification, achieving up to **95% accuracy**.
- Applied data augmentation, gradient clipping, and learning rate scheduling to handle class imbalance and improve generalization across 4 disease stages.

DroneX — UAV Object Detection System -GitHub *Sep 2022 – Jun 2023*

- Led AI and software team for UAV object detection; implemented **YOLO v4 Tiny** for real-time drone detection with field-collected training data. Secured **1st place** at the International UAV Competition.
- Programmed flight controllers using DroneKit; ran simulation and tuning in Gazebo on Linux; improved detection accuracy through iterative field testing.