

## Summary

- Machine Learning Engineer with 5+ years of research and industry experience spanning LLM applications, agentic systems, developer-facing AI software, generative models, and medical imaging. Experienced in building and evaluating LLM-powered systems, designing reliable agent workflows, and developing ML infrastructure and tooling for production environments.

## Skills

- LLMs & Agentic AI:** Agent Workflows, Guardrails, Agent Evaluation, Observability, Memory, RAG, Prompt and Context Management
- Machine Learning/AI:** Diffusion Models (Stable Diffusion, ControlNet), Variational Autoencoders (VAE), GANs, CNNs, U-Net, Self-Supervised Learning, Foundation Models
- Frameworks/Libraries:** PyTorch, Hugging Face, MONAI, TensorFlow, scikit-learn, scikit-image
- Software Engineering:** Developer-Facing Python Libraries, API Design, Package Development, Unit Testing, CI/CD, Git, Technical Documentation
- MLOps/Infrastructure:** AWS (SageMaker, EC2, S3), Docker, MLflow, ZenML, DVC, Slurm (HPC)
- Programming/Tools:** Python, C++, SQL, Bash/Linux, MATLAB

## Experiences

### Railtown AI – Machine Learning Engineer

Vancouver, Canada | Jan 2026 – Present

- Developed and maintained [Railtracks](#), a Python framework for building LLM-powered agents and applications, gaining hands-on experience with API design, abstractions, testing, documentation, and Python library development.
- Contributed to core abstractions for **agentic systems**, including LLM-powered agents, tool calling, context management, and integrations with external services and model providers.
- Worked on **LLM guardrails and reliability mechanisms** for controlling model inputs and outputs, including policy enforcement, content transformation, sensitive-data handling, and failure behavior.
- Developed **agent evaluation workflows**, including numerical and categorical metrics, tool-use evaluation, LLM-based evaluators, and LLM-as-a-judge approaches.
- Worked with **agent observability and tracing** to inspect agent runs, diagnose failures, and analyze model and tool interactions.
- Collaborated on the design of **developer-facing APIs and documentation**, translating complex LLM and agent concepts into reusable abstractions for production Python applications.

### Prenuvo – Machine Learning Intern

Vancouver, Canada | Sep 2024 – Dec 2025

#### Aortic Vessel Measurements and Aneurysm Detection – IRB-approved project

- Achieved **>90% Dice** for aortic section segmentation, **>90% sensitivity/specificity** for aneurysm detection, and **<0.5mm mean absolute error** for diameter measurement compared to radiology ground truth, meeting project performance targets.
- Collaborated with the clinical team** on data curation and annotation, ensuring high-quality inputs for algorithm validation.
- Integrated project into company's new AI monorepo**; worked with ML Engineering team to define repository structure, environments, and unit/performance tests, reducing integration and debugging time.
- Set up ZenML pipelines and orchestration scripts** for large-scale inference on AWS SageMaker, enabling scalable validation.
- Contributed to regulatory process** by drafting technical documentation and establishing data/model versioning, supporting successful IRB approval.

#### Heart Age Estimation – Feasibility Study

- Built nnU-Net segmentation model** on T2-weighted MRI, achieving **>90% Dice** for heart segmentation.
- Trained CNN-based regression network** for heart-age prediction with  $MAE \approx 4$  years, demonstrating the feasibility of organ age estimation.

- Transferred methods and data workflows (model training, inference, data selection) to the project lead, enabling continuation and scaling of the broader organ-age initiative.

#### Latent Diffusion Models for Brain Vessel Generation

- Trained Variational Autoencoder (VAE) on brain TOF MRI to learn compact image embeddings (SSIM > 0.85).
- Trained latent diffusion model for high-quality vessel image synthesis, enabling augmentation and diversity in the company's neurovascular datasets.
- Developed conditional diffusion (ControlNet) to guide generation using vessel segmentations, targeting anatomically precise image synthesis for downstream ML tasks.
- Integrated MLflow into diffusion model training pipeline, enabling systematic tracking of hyperparameters, metrics, and ablation studies for improved reproducibility.

### UBC – Research Assistant

Vancouver, Canada | Jan 2023 – Aug 2025 | Supervisor: Ilker Hacihaliloglu

- Designed and trained ControlNet/ControlNet++ diffusion models for synthetic echocardiogram generation, improving anatomical fidelity and diversity over GAN baselines.
- Augmented U-Net and Attention U-Net segmentation with diffusion-based synthetic data, achieving higher Dice scores and lower HD95 compared to real-only training.
- Introduced CMMD as a novel evaluation metric for synthetic data, offering strong alignment with segmentation outcomes.
- Optimized large-scale training on UBC ARC Sockeye cluster (Slurm, multi-GPU, mixed precision), reducing training runtime through efficient I/O and data loading.
- Published results at MICCAI 2025 and on arXiv (2024).

### Nabzgroup – Machine Learning Intern

Tehran, Iran | Jun 2021 – Sep 2021

- Implemented ECG R-peak localization and heartbeat classification algorithms for a handheld ECG recorder, combining signal processing with CNN and LSTM models.
- Applied wavelet-based preprocessing and post-processing pipelines to improve robustness of heartbeat detection under noise.
- Delivered end-to-end prototype achieving strong performance on benchmark datasets (MIT-BIH and in-house dataset), contributing to product development.

### Sharif University – Research Assistant

Tehran, Iran | Sep 2020 – Jun 2021 | Supervisor: M.B. Shamsollahi

- Implemented CNN-based classifiers for EEG motor imagery classification, achieving higher accuracy than classical EEG signal processing pipelines.
- Conducted experiments with PyTorch and MNE-Python, demonstrating the feasibility of deep learning for brain-computer interface (BCI) tasks.
- Published reproducible code and results on [Github](#), supporting transparency and future use in academic projects.

## Education

- M.A.Sc. in Biomedical Engineering, UBC

Vancouver, Canada | Jan 2023 – Aug 2025

GPA: 89.3/100

Thesis: Synthetic Echocardiogram Generation with Conditional Diffusion Models for Segmentation Tasks

[\[link\]](#)

- B.Sc. in Electrical Engineering, Sharif University

Tehran, Iran | Sep 2017 – Aug 2022

Specialization: Biomedical Engineering | Minor: Computer Science

GPA: 17.20/20

## Teaching & Mentorship

### UBC – Teaching Assistant

Vancouver, Canada | Sep 2023 – Aug 2025

- Supported large undergraduate classes (>200 students) in *Introduction to Data Science*, *Statistical Inference*, and *Elementary Statistics for Applications*.
- Led weekly tutorials covering probability/statistics concepts and R programming.
- Supervised and advised students on final projects, provided feedback on methodology and implementation.
- Held office hours and graded assignments/exams to support student learning.

## Publications

- Milad Yazdani, Yasamin Medghalchi, Pooria Ashrafian, Ilker Hacihaliloglu, Dena Shahriar. "Flow Matching for Medical Image Synthesis: Bridging the Gap Between Speed and Quality". *MICCAI 2025*. [\[link\]](#)[\[github\]](#)
- Pooria Ashrafian, Milad Yazdani, Moein Heidari, Dena Shahriari, Ilker Hacihaliloglu. "Vision-Language Synthetic Data Enhances Echocardiography Downstream Tasks". *Published on ArXiv*. [\[link\]](#)[\[github\]](#)