

Cristian Lazo Quispe

✉ clazoq@uni.pe | 📞 +51 970327308 | 🌐 Github | 💼 LinkedIn | 🎓 Google Scholar | 🌐 Portfolio

Computer vision researcher focused on **generative modeling and unsupervised anomaly detection for brain MRI**. Currently developing an age-conditioned flow matching framework for continuous brain-age interpolation (target: MICCAI 2027), mentored by Juan Miguel Valverde (DTU) through RISE-MICCAI; previously brain-MRI anomaly detection with Liam Chalcraft (UCL, Fatima Fellowship) and at Kyoto University. Two first-author challenge papers at **MICCAI 2025**; 1st place at the CVPR 2026 MSLR Challenge. 5+ years of production ML engineering, and committed to healthcare AI for low-resource settings - from a COVID-19 mechanical ventilator (OXIPUL) to teaching physicians clinical Big Data & AI.

Education

Kyoto University

Kyoto, Japan

Research Student in Intelligence Science and Technology

Sep 2024 - Sep 2025

Research focus on unsupervised anomaly detection in brain MRI using variational autoencoders and diffusion models.

Supervised by PhD. Roberto Espinoza Chamorro.

BREIT and the MIT Institute for Data, Systems, and Society (IDSS)

Lima, Peru

Advanced MicroMasters Program in Statistics and Data Science (**BREIT-MIT Scholarship**) Mar 2020 – Dec 2021

Graduate-level training in statistical learning, probability theory, and machine learning fundamentals.

Universidad Nacional de Ingenieria (UNI)

Lima, Peru

BSc in Mechatronic Engineering | Graduated with honors (top of class) (**Pronabec Government Full Scholarship**)

Mar 2015 – Dec 2019

Thesis: Scalable Face Recognition using Deep Learning and Vector Databases. *Supervised by Prof. Ricardo Rodríguez.*

Research Experience & Mentorship

RISE-MICCAI Research Mentorship – Research Mentee (🌐 web)

Remote, Sep 2025 – Present

Supervised by Postdoc. Juan Miguel Valverde, DTU Compute-Technical University of Denmark

- Working on **generative modeling for brain MRI** with applications in neurodegenerative disease analysis, including age-conditioned reconstruction frameworks.

Sign Language Recognition Lab – Research Volunteer (🌐 dict) (🌐 lab)

Remote, Jun 2022 – Sep 2025

AI4Good Lab, supervised by Professor Gissella Bejarano, Marist College

- Co-developed the **1st bilingual Peruvian Sign Language video dictionary** (text + video retrieval); conducted multilingual iconicity analysis using CLIP-based multimodal models. Co-authored publications at ICMV 2025, ISCMV 2025, and LatinX in AI workshops (CVPR 2025, NeurIPS 2022–2023).

Fatima Fellowship – Research Volunteer (🌐)

Remote, Jun 2023 – Dec 2023

Supervised by PhD Liam Chalcraft, University College London (UCL)

- Explored **unsupervised anomaly detection in brain MRI**, implementing Autoencoders, Variational Autoencoders (VAE), and β -VAE for synthetic data generation.
- Built baseline models for **anomaly localization and reconstruction-based scoring**, contributing to reproducible pipelines in PyTorch.

Achievements

🏆 1st place - CVPR 2026 MSLR Challenge, Track 2

2026

Multimodal Sign Language Recognition (0.931 Top-1 accuracy). Poster presented in person, Denver, CO.

🏆 1st place - UTEC Diploma in Big Data & ML against COVID-19

2021

Detection and severity classification of COVID-19 in postero-anterior chest X-rays using deep learning.

- | | | | |
|---|------|--|------|
| 1. Top 7% out of 1k+ participants WiDS Datathon | 2025 | 5. Top 3% out of 10k+ participants WiDS Datathon | 2022 |
| 2. Top 1 Datathon Ferreyros Forecasting | 2024 | 6. Top 1 Datathon Ricardo Palma University | 2019 |
| 3. Top 2 Datathon Marine Audio Recognition | 2022 | 7. Top 1 Datathon BNP Paribas & Coursera | 2021 |
| 4. Top 1 Datathon Entel Telecom Forecasting | 2022 | 8. Top 6 Programming Contest ACM-ICPC Regional | 2017 |

Publications

Under review / in progress

- Flow Matching with Contrastive Learning for Continuous Brain Age Interpolation in MRI
[Cristian Lazo](#), Juan Miguel Valverde **In progress - target MICCAI 2027**
Reconstruction is strong; the current frontier is age-conditioned trajectory generation, where error grows with distance to the target age - I am iterating on contrastive regularization and trajectory-straightening to address it.

Challenge tracks @ MICCAI 2025

- Robust Multi-Label Classification of MRI Artifacts in Low-Field Neonatal Brain Imaging via View-Conditional Dual-Task Learning ( 
[Cristian Lazo](#), Roberto Espinoza **LISA Challenge @ MICCAI 2025** (*Springer LNCS, Scopus Q2*)
- Cascaded Coarse-to-Fine Segmentation for Hippocampal Structures in Low-Field Neonatal MRI (
[Cristian Lazo](#), Roberto Espinoza **RISE Challenge @ MICCAI 2025**

Challenge tracks @ CVPR 2026 Workshops

- AI4Good at SignEval 2026: Mix, Measure, Merge with Spectro-Temporal Augmentation and Multi-Model Fusion for Isolated Radar-Based Italian Sign Language Recognition (
[Cristian Lazo](#), Gissella Bejarano **1st place - MSLR Challenge @ CVPR 2026** (*CVPRW proceedings*)
- Robust Deepfake Detection, NTIRE 2026 Challenge: Report (
B. Hopf, R. Timofte, *et al.* (incl. [Cristian Lazo](#)) **NTIRE Challenge @ CVPR 2026** (*CVPRW proceedings*)

Peer-reviewed (conference / journal)

- WholeBodyPose: A Unified End-to-End Framework for SLR and Pose-Based Training Data (
[Cristian Lazo](#), Renato Castro-Cruz, Gissella Bejarano **ICMV 2025**
- Real-Time Face Tracking and Vector Databases for Scalable Face Recognition (
[Cristian Lazo](#), Renato Castro-Cruz, Ricardo R. Rodríguez Bustinza **ISCMV 2025**
- Automatic 3D Teeth Segmentation for STL Teeth Models (
[Cristian Lazo](#) **IHIP 2019**
- A Brief Survey on DL-Based Methods for Lung Cancer Classification (
Mario G. Borja Borja, Roger Huauya, [Cristian Lazo](#) **IEEE CHILECON 2019**
- Static Hand Gesture Recognition for Peruvian Sign Language (
[Cristian Lazo](#), Zaid Sanchez, Christian Carlos del Carpio **BTSym 2018**

Workshops & non-archival - LatinX in AI

- SignERT: An Efficient Pose-Only Transformer for ISLR ( Demo)
[Cristian Lazo](#), Mauricio Salazar, Gissella Bejarano **LatinX in AI @ CVPR 2025** (*extended abstract*)
- Enhancing Anomaly Detection with Spatial Transformer Networks (
Renato Castro-Cruz, [Cristian Lazo](#) **LatinX in AI @ NeurIPS 2023**
- Impact of Pose Estimation Models for Landmark-based SLR (
[Cristian Lazo](#), Joe Huamani, Gissella Bejarano, Pablo Rivas, Tomas Cerny **LatinX in AI @ NeurIPS 2022**
- Segmentation of Skin Lesions and Their Attributes Using GANs (
[Cristian Lazo](#) **LatinX in AI @ NeurIPS 2019**

Selected Research Projects & Demos

- **Generative Modeling – Diffusion & Flow Matching Notes** (  Demo) ( Flow,  DDPM) 2025
Extensive mathematical derivations and PyTorch implementations of DDPM/Flow Matching, with intuitive explanations of ELBO, variational inference, Neural Flows, and ODE.
- **Multimodal Image–Text Matching** ( 2025
Exploration of image–text alignment using CLIP, DeepSeek, and hybrid architectures (ResNet–BERT, DeepSeek–CLIP). Trained and compared multiple multimodal encoders for semantic retrieval.
- **Spontaneous Behavior & Visual Cortex Dynamics** ( 2022
Analysis of large-scale neural recordings from mouse visual cortex (Stringer et al.), uncovering correlations between spontaneous behavioral variables and population-level neural activity.

- **OXIPUL Mechanical Ventilator - Software & Embedded Systems Lead** ( [press](#)) 2020
Led the real-time control and monitoring software of a COVID-19 mechanical ventilator (pressure/temperature telemetry with live visualization, communication protocols, and motion-control algorithms solved numerically in real time) at the UNI Faculty of Mechanical Engineering, funded by Compañía Minera Las Bambas. Fully functional prototype presented to and endorsed by Peru's Minister of Health, Dra. Pilar Mazzetti; covered by national press.

Teaching

- Universidad Científica del Sur** Lima, Peru
Lecturer - Health Data Science (DICDASA) Sep 2025 – Present
 Teach two courses of the Postgraduate Diploma to physicians, clinical managers and hospital administrators, without assuming programming: *Principios de Big Data* (Big Data foundations, distributed architectures, responsible AI) and *ML & AI Methods in Health* (supervised, unsupervised and transfer learning; evaluation and clinical failure modes).
- BREIT & MIT IDSS (Aporta-MIT partnership)** ( [web](#)) Remote, Peru
Teaching Assistant - Probability, Statistics & Machine Learning May 2022 – Present
 Co-taught the MIT MicroMasters in Statistics & Data Science for cohorts in Peru (Machine Learning, Probability, Data Analysis in Social Science); mentored end-to-end ML projects in PyTorch and HuggingFace. Authored and published all **13 review sessions** of MIT 14.310x: slides, theory notes, a five-level problem bank and full solutions.
- 4Geeks Academy** Remote, Uruguay
Teaching Assistant - Machine Learning May–Sep 2022
 Supported training of Latin American professionals in probability, statistics, data analytics, and applied machine learning.
- Data Science Women (Data Science Research Perú)** Remote, Peru
ML Teaching Assistant & Project Lead Apr 2020 – Apr 2022
 Mentored and led collaborative ML projects, guiding diverse teams and supporting women in tech.
- CTIC - National University of Engineering (UNI)** Lima, Peru
Volunteer Instructor - Machine Learning & Computer Vision Jun 2018 – Dec 2019
 Taught ML/DL fundamentals and computer vision projects (C++, Python) to undergraduate students. Developed prototypes for anomaly detection, license plate recognition, and object tracking.

Industry Experience

- Marvik** ( [web](#)) Remote, Uruguay
Senior Machine Learning Engineer Nov 2025 – Present
 AI/ML consulting for global clients including Fortune 500 companies; designing and deploying end-to-end ML solutions in computer vision, NLP, and multimodal AI. (Marvik is an NVIDIA Inception Programme member.)
- Nisum Latam** Remote, Chile & USA
AI Technical Lead Jan 2023 – Sep 2024
 Led enterprise LLM, multi-modal, and computer vision projects (document understanding, audio diarization, named entity recognition, code migration, and visual QA automation) and deployed scalable pipelines with Docker, Kubernetes, and CI/CD workflows.
- Rimac Seguros** Lima, Peru
Machine Learning Engineer May 2021 – Jan 2023
 Built production-grade ML pipelines (tabular, images, text, and audio data) on GCP Vertex AI; led deployed human-in-the-loop annotation for computer vision models that reduced manual inspection time significantly.
- MDP Consulting** Lima, Peru
Machine Learning Engineer Oct 2019 – Jan 2021
 Developed real-time facial recognition and license plate recognition systems for edge deployment (Jetson TX2), integrating cloud-based APIs and CV models into operational workflows. Built a masked/unmasked face recognition and access-control system for COVID-19 safety protocols, featured on national TV news.

Skills

Computer Vision & Biomedical Imaging: Medical Image Analysis (MRI, X-Ray, CT-scans), 3D Reconstruction and Generation, Semantic Segmentation, Anomaly Detection, Pose Estimation.

Generative Models: Diffusion Models, Flow Matching (Rectified Flow, OT-CFM), VAEs, GANs.

Biomedical Imaging Toolkits: MONAI, 3D Slicer, MITK, MeshLab.

Multimodal ML: CLIP, Sign Language Transformers, Speech-to-Text, OCR + CV + LLM fusion.

Programming & Frameworks: PyTorch, OpenCV, Scikit-learn, MLFlow, Weights & Biases, CI/CD (GitHub Actions).

Cloud & Infrastructure: GCP, AWS, Databricks, Kubernetes, Docker, Remote HPC.

Distributed Training: Multi-node / multi-GPU training with PyTorch DDP and FSDP, SLURM, mixed precision, gradient checkpointing, sharded optimizer state; MLFlow, Weights & Biases and PyTorch Profiler for tracking and throughput analysis.

Embedded Systems for Edge AI: Jetson Nano, Triton, ONNX, Raspberry Pi 4, ESP32; Real-Time Pipelines.

Programming Languages: Python, C++, Matlab, C, Assembler.

Languages: Spanish (Native), English (Advanced), Japanese (Elementary).

Certifications

[0] DL4Sci Summer School – Lawrence Berkeley National Lab & UC Berkeley 2026
Deep Learning for Science; selected participant (competitive), Jul 20–24. Foundation models, reasoning and agentic AI for science. Presented the poster *Age-Conditional Latent Flow Matching for Modeling Healthy Brain Ageing in MRI*.

[1] Oxford Machine Learning Summer School – Oxford MLx 2025
MLx Fundamentals, Generative AI, Health & Bio, Representation Learning.

[2] Large-Scale Training in the Modern World – Maven & Hugging Face 2025
Torch, SLURM, Parallelism, Large-Scale Training of Multimodal Models.

[3] How Diffusion Models Work – DeepLearning.AI 2025
Diffusion Models, PyTorch, Probabilistic Models, Computer Vision.

[4] NMA Computational Neuroscience Program – Neuromatch Academy 2022
Computational Neuroscience, Neural Coding, ML, Python, Brain Modeling.

[5] Capstone Exam – MicroMasters in Statistics and Data Science – MITx 2021
Integration of Machine Learning, Statistics, and Data Analysis.