Zohreh Mostaani

Machine Learning Engineer

mostaani.zohreh@gmail.com | Switzerland | zmostaani.github.io
Zohreh-mostaani | t2VMpU8AAAAJ | Swiss C Permit, Driving License



SUMMARY

Machine Learning Engineer with a PhD (EPFL) and 7+ years building deep learning and time-series ML pipelines end to end: data acquisition, modeling, and rigorous evaluation. Deep ML experience across a broad range of domains: speech, physiological and health signals, radar, and biometrics. Strength in rigorous ML evaluation across conditions, including a recent automated evaluation framework for generative speech, validated against human listening studies. Experienced with self-supervised speech representations (wav2vec 2.0, HuBERT), PyTorch, and containerized ML tooling (Docker, REST APIs, automated testing).

SKILLS

Languages & Core: Python (NumPy, Pandas, SciPy), Shell, MATLAB; Git, Docker

ML & Deep Learning: PyTorch, TensorFlow, scikit-learn; CNNs, RNNs, transformers/attention; time-series modeling; self-supervised speech representations (wav2vec 2.0, HuBERT)

ML Engineering & Systems: Containerized microservices, REST APIs (FastAPI/Flask); automated testing; environment management (conda/mamba, Poetry, uv); LLM-assisted development (Claude Code, opencode)

Evaluation: Automated evaluation frameworks; benchmarking across conditions; failure-mode analysis; human-correlation / listening studies (MOS); ablation

Cloud & Data: AWS S3, Azure ML; dataset design & collection, labeling protocols, data pipelines/loaders; signal processing; visualization (Matplotlib/Seaborn/Plotly)

Domains: Speech & audio; physiological & health signals (breathing, heart-rate, hypoglycemia, COVID-19 from speech); radar vital-sign monitoring; face PAD / biometrics

Spoken languages: English (fluent), French (intermediate), Persian (native)

EXPERIENCE

Research Engineer
Jan 2026 – Jun 2026

AGIGO AG
Zürich, Switzerland

Generative speech evaluation for an agentic, LLM-driven conversational AI platform.

- Designed and built an automated evaluation framework that scored 9 TTS voices across 15 domains end to end, taking text as input and producing a single quality score (intelligibility, audio quality, speaker similarity, synthesis success rate); reduced reliance on manual listening tests and validated against a human ranking study, matching it on the top-3 and bottom-2 voices.
- Built a diagnostic tool that localized TTS errors via an ASR round-trip and classified mismatches by severity, letting developers trace each failure to the TTS, ASR, or normalization stage.
- Packaged the framework as a Dockerized microservice with HTTP endpoints and integrated it into the team's model-evaluation workflow.
- Set up the project's automated test environment in line with company-wide infrastructure standards.
- Advised three research interns working on TTS.

Machine Learning Engineer (BNF Project)
Jul 2025 – Oct 2025

Sleepiz AG
Zürich, Switzerland

- Designed and ran data-collection protocols for radar-based, non-contact human detection and vital-sign monitoring, and collected the data used for algorithm development.
- Developed a distance-bin detection method (signal processing and filtering) that reduced MSE by ~50% versus the default fixed-bin baseline.
- Fed the detected bins into an existing ML model for vital-sign estimation, cutting downstream errors versus the fixed-bin baseline by ~31% for breathing rate and ~87% for heart rate.

Research Assistant - Speech and Audio Processing
Feb 2020 – Mar 2025

Idiap Research Institute
Martigny, Switzerland

Integrated physiological signals in speech processing using machine and deep learning

- Developed CNN models to estimate breathing signals directly from raw speech waveforms (TensorFlow), reducing input lag by 25% versus spectral-based baselines; the embeddings from these models became a reusable component across downstream health and paralinguistic tasks.

- Built hypoglycemia-detection models from speech and ECG for patients with type 1 diabetes, combining CNN and self-supervised embeddings with handcrafted acoustic features; reached up to 96% AUC and validated results across two clinical datasets, contributing to a peer-reviewed publication (Diabetes Care, 2026).
- Applied the same representation-learning approach across a range of speech-based health and paralinguistic tasks, including heart-rate estimation (~0.75 correlation, improving ~20% with longer input windows), emotion recognition, COVID-19 detection, and detecting breathing patterns in synthetic speech (99% AUC on TTS-generated audio).
- Designed and ran rigorous evaluations that surfaced strong per-individual variation.
- Designed acquisition protocols and collected two multimodal speech-and-physiological databases from 54 participants total (22 in clinical settings, in collaboration with medical doctors); assessed the accuracy limits of a wearable sensor.
- Mentored two Master's students; TA for Master's ML and Speech Processing courses.

Research and Development Engineer - Biometrics Security and Privacy Martigny, Switzerland
Idiap Research Institute Apr 2019 – Jan 2020

- Restructured and rewrote the documentation for the BEAT ML-experimentation framework (Python, Django, Sphinx), making an outdated, unusable platform accessible again; performed QA testing across experiment execution and usability, filed detailed reproducible bug reports, and fixed bugs.
- Collected a face-video database with a makeup-based aging protocol, showing it could lower a face-recognition system's median similarity score by 14% demonstrating an effective presentation attack; co-authored the resulting paper proposing a CNN-based detection method.

Research Internships Jun 2016 – Mar 2019
Idiap Research Institute Martigny, Switzerland

- Collected two large face-biometric datasets from 100+ participants, including acquisition-protocol design and generation of presentation attacks.
- Implemented a CNN for gaze estimation (Theano); integrated an ASR system into a web framework using Docker and a REST API.

Research and Teaching Assistant Feb 2012 – Jan 2015
Özyeğin University Istanbul, Türkiye

- Research: improved signal-to-noise ratio by 50% in free-space optical communication using multi-relay MIMO systems (MATLAB, Python); studied outage performance under varying turbulence.
- Teaching assistant for Bachelor's and Master's courses in physics, mathematics, linear algebra, and wireless communications.

EDUCATION

PhD, Electrical and Electronics Engineering, **EPFL** Mar 2025
M.Sc., Electrical and Electronics Engineering, **Özyeğin University** Jan 2015
B.Sc., Electrical Engineering, **University of Tehran** Sep 2011

PUBLICATIONS

15+ peer-reviewed papers in venues including Neural Networks, ICASSP, INTERSPEECH, and Diabetes Care (~700 citations, h-index 9 Google Scholar). Reviewer for IEEE/ACM Transactions on Audio, Speech, and Language Processing. Selected:

- Deep Learning Architectures for Estimating Breathing Signal and Respiratory Parameters from Speech – Neural Networks, 2021
- Listening to Hypoglycemia: Voice as a Biomarker for Detection of a Medical Emergency Using Machine Learning – Diabetes Care, 2026
- Biometric Face Presentation Attack Detection with Multi-Channel CNN – IEEE T-IFS, 2020

SELECTED PROJECTS

- ***AI for Children's Reading Assessment*** *(Lemanic Life Science Hackathon, EPFL): built a working prototype to identify target words and judge pronunciation correctness from a child's speech (team of 10, three days).*
- **RCS *Reduction of a Patch Antenna*** *(B.Sc. thesis, University of Tehran): reduced radar cross-section using Particle Swarm Optimization (MATLAB, FEKO).*

AWARDS

Full PhD scholarship (EPFL); full merit scholarships, Özyeğin University and University of Tehran. Ranked 206th of ~400,000 in Iran's nationwide university entrance exam (2007).