Fernando Pérez-García (he/him)

fepegar@gmail.com | LinkedIn | GitHub | Scholar | ORCID

Profile

I am an engineer with a PhD in Medical Imaging and a decade of experience applying cutting-edge multimodal AI for healthcare to solve real-world problems and supporting biomedical research. I strongly believe in the importance of meaningful translational research and have developed open-source machine learning software actively used by both clinicians and researchers. I enjoy being halfway between research and engineering and have developed strong communication skills to connect ideas from both worlds.

Experience

Senior Researcher and Team Lead (*Microsoft Health Futures, Cambridge, UK*) *Feb 2022 – Present*

- Focus on Multimodal AI for healthcare: self-supervised learning (e.g. RAD-DINO, Nature Machine Intelligence) and generative diffusion models for medical imaging (e.g. RadEdit, ECCV)
- Collaborated with the NHS to develop a deep learning model for hippocampal segmentation from brain MRI scans
- Collaborated with Mayo Clinic to implement our multimodal LLM, MAIRA, in their clinical workflows

Research Software Engineer (freelance) *Jan 2019 – Sep 2021*

- Paris Brain Institute (ICM): Developed tools to build a 3D atlas of the macaque brainstem (histology & MRI)
- Arterys – Medical Imaging Cloud AI: Deployed deep learning models for medical images using Docker
- Medical Augmented Intelligence (MAI): Developed virtual reality and medical imaging software on 3D Slicer

Research Software Engineer (*Paris Brain Institute (ICM) – CENIR, Paris, France*) *Aug 2015 – July 2017*

- Created 3D atlases of human and monkey brainstems based on immunohistochemical and MRI data
- Collaborated with Medtronic on a product for planning and assessment of stereotactic surgery

Research Software Engineer (intern) (*Center for Biomedical Technology (CTB – UPM), Madrid, Spain*) *Nov 2014 – Feb 2015*

- Collaborated with NASAL Inc. on a DICOM reading/writing software tool for planning of sinus surgery

Research Software Engineer (intern) (*Paris Brain Institute (ICM) – CENIR, Paris, France*) *Mar 2014 – Sep 2014*

- Developed EpiLoc, an automated pipeline for planning and assessment of electrode implantation surgery in epilepsy

Education

PhD in Medical Imaging — *University College London & King's College London* *Sep 2018 – Nov 2022*

- Lead the development of TorchIO, a library for medical images and deep learning (in PyTorch ecosystem)
- Coordinated a collaboration between UCL and UCLH to publish EPISURG, a large MRI dataset
- Wrote and released Resseg, a tool for quantitative analysis of resective neurosurgery using deep learning

- Collaborated with neurophysiologists at NHNN to classify seizures from video-telemetry using deep learning
- Joined the Medical Open Network for AI (MONAI) Transformations Working Group as an advisor

MRes Medical Imaging — *University College London* *Sep 2017 – Sep 2018*

MEng Biomedical Engineering — *Universidad Politécnica de Madrid* *Sep 2014 – July 2015*

BEng Sound and Image Engineering — *Universidad Politécnica de Madrid* *Sep 2009 – July 2014*
- Erasmus year at Beuth Hochschule für Technik Berlin, Germany (2011/12)

BEng Industrial Electronics Engineering — *Universidad Pontificia Comillas — ICAI* *Sep 2007 – July 2010*

Technical Skills

Languages: Python (since 2012), MATLAB (10 years), C (5 years), JavaScript (a few years), shell (since 2015)
Tools: Azure Cloud, PyTorch, Docker, Git[Hub], uv, pytest, tox, VS Code, TorchIO, MONAI, ITK, VTK, 3D Slicer

Languages

Spanish: native; English (IELTS – 2017) and French: full professional proficiency; German (Goethe-Zertifikat B2 – 2012) and Italian: intermediate; Chinese (HSK2 – 2013) and Arabic: beginner.

Selected Publications

- Fernando Pérez-García, Harshita Sharma, Sam Bond-Taylor, Kenza Bouzid, Valentina Salvatelli, Maximilian Ilse, Shruthi Bannur, Daniel C Castro, Anton Schwaighofer, Matthew P Lungren, Maria Teodora Wetscherek, Noel Codella, Stephanie L Hyland, Javier Alvarez-Valle, and Ozan Oktay. **Exploring scalable medical image encoders beyond text supervision.** *Nature Machine Intelligence* 2025.
- Shruthi Bannur, Kenza Bouzid, Daniel C Castro, Anton Schwaighofer, Sam Bond-Taylor, Maximilian Ilse, Fernando Pérez-García, Valentina Salvatelli, Harshita Sharma, Felix Meissen, Mercy Ranjit, Shaury Srivastav, Julia Gong, Fabian Falck, Ozan Oktay, Anja Thieme, Matthew P Lungren, Maria Teodora Wetscherek, Javier Alvarez-Valle, and Stephanie L Hyland. **MAIRA-2: Grounded Radiology Report Generation.** *arXiv preprint arXiv:2406.04449* 2024.
- Fernando Pérez-García, Sam Bond-Taylor, Pedro P Sanchez, Boris van Breugel, Daniel C Castro, Harshita Sharma, Valentina Salvatelli, Maria TA Wetscherek, Hannah Richardson, Matthew P Lungren, Aditya Nori, Javier Alvarez-Valle, Ozan Oktay, and Maximilian Ilse. **RadEdit: stress-testing biomedical vision models via diffusion image editing.** *European Conference on Computer Vision (ECCV)* 2024.
- Shruthi Bannur, Stephanie Hyland, Qianchu Liu, Fernando Perez-Garcia, Maximilian Ilse, Daniel C Castro, Benedikt Boecking, Harshita Sharma, Kenza Bouzid, Anja Thieme, Anton Schwaighofer, Maria Wetscherek, Matthew P Lungren, Aditya Nori, Javier Alvarez-Valle, and Ozan Oktay. **Learning to exploit temporal structure for biomedical vision-language processing.** *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 2023.
- Fernando Pérez-García, Rachel Sparks, and Sebastien Ourselin. **TorchIO: a Python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning.** *Computer Methods and Programs in Biomedicine* 2021.
- Andres Diaz-Pinto, Sachidanand Alle, Vishwesh Nath, Yucheng Tang, Alvin Ihsani, Muhammad Asad, Fernando Pérez-García, Pritesh Mehta, Wenqi Li, Mona Flores, Holger R Roth, Tom Vercauteren, Daguang Xu, Prerna Dogra, Sebastien Ourselin, Andrew Feng, and M Jorge Cardoso. **Monai label: A framework for ai-assisted interactive labeling of 3d medical images.** *Medical Image Analysis* 2024.