Rafael Felix, Ph.D.

Research Scientist

Researcher scaling AI models to real-time applications serving 1B+ daily users with p99 <2s latency across billions of requests. Focused on multimodal AI for Trust & Safety and Robotics. Published in top conferences on multimodal, generative, zero-shot learning, and safety-critical systems. Experienced in team leadership, mentoring, and driving organizational impact.

Experience

Research Scientist in Multimodal AI

Company: TikTok, in Remote, from Apr 2024 to Aug 2025

- **Minor Safety Business:** Improved age estimation in short videos, boosting benchmark accuracy by +12% and raising safety precision from 30%→40% across 1B daily moderation tasks. Partnered with product and policy teams to align models with global safety standards.
- **Scalable Benchmarks & Compliance:** Built a 100K+ annotated dataset (expanding 5K/week), addressing data/model drift and enabling regulatory compliance. Designed workflows that streamlined efficiency and supported long-term cross-team use.
- **Real-Time Inference Optimization:** Achieved p99 <2s latency at 11K QPS via RPC-based inference, using CUDA, TensorRT, and ONNX quantization for scalable deployment in production pipelines.
- **Leadership & Impact:** Led team of 4 ICs, mentored researchers, and drove adoption across 50+ researchers through workshops and best practices. Collaborated cross-functionally with infra and moderation teams to integrate research into large-scale systems. Published at ICCV.

Research Scientist

Company: Amplified Intelligence, in Australia, from Apr 2022 to Mar 2024

- **Cross-Device Gaze Estimation:** Reduced point-of-gaze error from ~50mm to <30mm across mobile, desktop, and TV. Delivered business-aligned CNN/ViT fine-tuning while ensuring regulatory compliance.
- **4K Multimodal Engagement Model:** Built the first model combining gaze, pose, and epipolar geometry for 3D projection, achieving 84% F1 in pedestrian ad engagement; set a new technical and business benchmark.
- **Scalable Inference at 4K:** Designed pipeline supporting 8K QPS and 360MB/s video throughput via CPU auto-scaling with TFLite/ONNX, enabling efficient deployment at scale.
- **Industry-Grade Benchmarks & Compliance:** Established a dataset of 5K+ subjects to address model drift, shape GDPR/COPPA compliance, and drive company-wide standards with reusable pipelines.
- **Leadership & Impact**: Influenced technical direction with C-suite; led a team of 4 ICs; mentored 20+ engineers via 1:1s and workshops, scaling ML adoption and operational efficiency across the org. Published at ECCV/WACV.

Research Scientist

Company: Australia Inst Machine Learning, in Australia, from Jan 2020 to Mar 2024

- **Vision-Language Systems for Surveillance:** Improved VQA accuracy by +5% on satellite and illegal fishing datasets. Built a chatbot-style VQA system that reduced query resolution time 30×, demonstrating operational gains for government and industry partners.
- **Generative AI & Deepfake Detection:** Benchmarks improved from 0.41 → 0.98, with models integrated into multiple research programs and collaborations. Developed GANs and diffusion models for style transfer, fake detection, and dataset augmentation to improve zero-shot learning.
- **Research Strategy & Funding:** Secured grants and authored high-impact papers, shaping academic and industry discussions on LLMs for vision-language tasks. Influenced a research group of 200+ scientists and engineers through workshops and mentorship.

Research Scientist

Company: Australia Centre Robotic Vision, in Australia, from June 2016 to Nov 2019

- **Zero-Shot Learning & Generative AI:** Published 5 papers pioneering the use of generative models for data augmentation, enabling Zero-Shot Learning systems to recognize novel concepts beyond training data. Advanced GAN-based pipelines for unseen class image generation, pushing state-of-the-art in vision-language models for robotics.
- **Research adoption:** Work adopted by 400+ researchers, advancing vision-language models through highly cited top-tier papers and open-source code (57 stars, 22 forks) that enabled reproducibility and community extension.

Education

The University of Adelaide, Australia
Doctorate, Artificial Intelligence and Robotics, from Jun 2016 to Nov 2019
Mackenzie Presbyterian University, Brazil
Masters, Artificial Intelligence and Robotics, from Jan 2013 to Jun 2015
MC State University, Brazil
Bachelor, Computer and Information Science, from Jan 2007 to Jun 2011

Skills

Python, PyTorch, TensorFlow, SQL, Docker, Kubernetes, C++, Spark
Transformers, LLMs, CNNs, Diffusion, GANs, ViTs, BERT, CLIP, LLaVA, LLaMA, Qwen
Pre-training, Fine-tuning, Transfer Learning, RLHF, RNN, LSTM, Zero-Shot Learning, Gaze Estimation
English, Portuguese, Spanish, Japanese

Websites

rafafelix.info
github.com/rfelixmg
linkedin.com/in/rafafelixphd
scholar.google.com.au/citations?user=nijDcmQAAAAJ

Research papers

G,A.,N,C.,Felix,R.,Do,T.-T.,&G,G. AEON: Adaptive Estimation of Instance-Dependent In-Distribution and Out-of-Distribution Label Noise for Robust Learning (2025)

G,A.,N,C.,Felix,R.,Do,T.-T.,&G,G. Instance-Dependent Noisy-Label Learning with Graphical Model Based Noise-Rate Estimation. ECCV(2024).

Hu, H., Chen, H., Felix, R., Abbasnejad, E., & Lingqiao, L. Progressive Feature Adjustment for Semi-supervised Learning from Pretrained Models. ICCV(2023).

G, A., N, C., Felix, R., Do, T.-T., & G, G. Instance-dependent noisy label learning via graphical modelling, WACV (2023)

G, A., Felix, R., S, J., A, E., & v. d. H., A. Colour palette generative adversarial networks. arxiv (2023).

Felix, R., S, J., A, E., & v. d. H., A. Cross-modal visual question answering for remote sensing data. DICTA(2021).

Felix, R., S, M., R, I., & G, G.. Augmentation network for generalised zero-shot learning. ACCV (2020)

Felix, R., H, B., S, M., & H, G. Gen. ZSL with domain classification in a joint semantic and visual space. DICTA(2019).

Felix, R., S, M., Reid, I., & G, G. Multimodal ensemble classification for generalised zero shot learning. arxiv (2019).

Felix, R., R, I., G, G. et al. Multi-modal cycle-consistent generalised zero-shot learning. ECCV (2018).

Felix, R., da S, L. A., & de C, L. N. Thresholding the courtesy amount of Brazilian bank checks using a local methodology. Springer ICPAAMAS (2015).