Tharen Candi

+39 3761273583 | Padova, Italy / Sydney, Australia | candi.tharen@gmail.com | GitHub | LinkedIn | Website

Italian/Australian dual citizen. MSc Computer Science (AI) student at the intersection of AI and systems: TinyML, embedded systems, computer vision, representation learning. Graduating April 2027, and seeking a graduate role.

Education

University of Padova, MSc Computer Science (AI) · GPA: 28.08/30 — Oct 2024 — Apr 2027
University of Sydney, BComp (Computer Science) · GPA: 82.2/100 (Distinction) — Feb 2020 — Dec 2023
EPFL, Exchange Program · GPA: 5.6/6.0 — Sep 2022 — Feb 2023

Research and Projects

MSc Thesis: Closed-Loop On-Device Perception for Tiny Robots, MIG, IAS-Lab — Sep 2026 — Present
- Integrating Rust TinyML inference (microflow-rs) into RIOT OS on an ESP32-S3 for closed-loop navigation.
- Contributing camera support to RIOT OS upstream: an SPI driver for the Arducam OV2640 (JPEG and raw).
- Developing and optimising TinyML perception models and the on-device inference pipeline.

The Limits of Cross-Lingual Transfer: SignCLIP on LIS, Research · Paper · GitHub — Feb 2026 — Jun 2026
- Matched SOTA on A3LIS-147 with a fine-tuned linear probe (R@1 80.2%).
- First SignCLIP study on Italian Sign Language: zero-shot transfer helps iconic signs, hurts non-iconic ones.
- Ablated six contrastive losses: GlobalNCE fine-tuning raised zero-shot R@1 from 3.7% to 76.2%.

Universal Sign Language Representation Learning, Independent Research — Aug 2026 — Present
- Targeting the negative cross-lingual transfer found in SignCLIP with language-universal sign embeddings.
- Designing a language-gated, JEPA-inspired pose model that predicts into a frozen text-embedding space.
- Early results show above-random transfer to sign languages unseen in training (ongoing).

3D Pinna Landmark Extraction, Huawei Munich Tech Arena — Aug 2026 — Sep 2026
- Predicted 85 ear landmarks from 3D head scans at ~1.2 mm mean error, for personalised spatial audio.
- Ranked 12th of 38 teams (top third) as an all-Master's team in a challenge that recommended PhD members.
- Ensembled DGCNN, HRNet and Point2SSM++ for 15 anchors; traced 70 contour points as shape-guided geodesics.

Multimodal Speech Emotion Recognition · Paper · GitHub — Jun 2026 — Jul 2026
- Fused speech (wav2vec 2.0) and facial (MediaPipe + XGBoost) models for 7-class emotion recognition.
- Calibrated late fusion improved accuracy by 5.3–7.7 points across models (best: 78.5%).
- Under matched probes, wav2vec 2.0 and emotion2vec showed no statistically reliable difference.

Hard Real-Time: Ada Ravenscar vs. Rust RTIC · Paper · GitHub — Jan 2026 — Feb 2026
- Ported a hard real-time benchmark from Ada Ravenscar to Rust RTIC on ARM Cortex-M4; modelled both in MAST.
- Measured RTIC cutting interrupt latency by 45% and scaling to 100 tasks, where Ada ran out of memory at 25.
- Showed RTIC's global interrupt masking adds non-preemptive blocking that schedulability analysis must model.

Glioblastoma Cell Segmentation Pipeline, Undergraduate Capstone · GitHub — Feb 2023 — Aug 2023
- Trained a DRAN CNN in TensorFlow to segment gigapixel whole-slide glioblastoma images (TCGA-GBM).
- Deployed it as an automated pipeline built on OpenCV, Celery and Redis.

Work Experience

Graduate Intern – Carbon Trading and Origination, Glencore — Jul 2024 — Sep 2024
- Automated forecasting ingestion, cleaning and ArcGIS workflows for the CER ACCU Register, saving 5 h per cycle.
- Built supply-side carbon market models and analytics in Power BI and Excel.
- Documented the automated workflows for technical and non-technical users so the team could maintain them.

Online Department Shift Leader, Coles, Clemton Park, Sydney — Aug 2021 — Aug 2022
- Led online order fulfilment for pickup and delivery, managing time-sensitive logistics and customer service.

Honours and Awards

- **Master's Study Grant**: Italian MFA (2024, 2025)
- **Dean's List**: University of Sydney (2023)
- **Global Mobility Scholarship**: USYD (2022)
- **HSC Top Achiever**: 6th in State for IPT (2019)

Skills

- **Languages**: Python, C/C++, Rust, Ada, Java, Scala
- **ML**: PyTorch, TensorFlow, Hugging Face, OpenCV
- **Systems**: RIOT OS, RTIC, ROS 2, Linux, CUDA, Git
- **Spoken**: English (native), Italian (B2), Greek (A1)