

Wei-Cheng Wang, Ph.D.

tienjiwang@gmail.com | linkedin.com/in/wang-wc | [Personal Website](#) | (+886) 919242749 | Taiwan

Summary

Independent Researcher with **10+ yrs research** experience on **unlabeled streaming data** (video, audio, and multimodal data), including extended **video** streams and **5000+ hours** of **uncurated audio** collected across multiple cities and years. Research bridges **lab to real-world deployment**, tackling **domain shift**, **source-free** and **unsupervised** settings, while providing solutions in **privacy-preserving ML** and **general-purpose** representation. Recognition includes 13 peer-reviewed **articles (150+ citations)** with **4 first-author journals** (2 Top 10%, 1 Top 25%) and **ICIP** (first-author, oral), **APICTA** award (team lead, representing Taiwan), and **Top 3%** master's graduate honor.

Skills

- **Domain Expertise:** Computer Vision, Machine Learning, Multimodal ML & VLMs, Quantization & Edge Deployment, Self-Supervised & Contrastive Learning, Anomaly Detection, Privacy-Preserving AI, Generative Models.
- **Tech Stack:** Python, PyTorch, TensorFlow, OpenCV, quantization and edge deployment toolchains (ONNX, AIMET, Qualcomm AI Hub), Docker, AWS.

Experience

Independent Researcher | Dec 2025 - Present

Python, PyTorch, ONNX, AIMET, TorchAO, Qualcomm AI Hub (QNN), AWS (ECS, S3), Docker

- **Cross-environment (GPU/NPU) evaluation of edge-cloud VLM quantization:**
Evaluated **vision encoder quantization** at the **on-edge** embedding, refittable downstream heads and the **on-cloud LLM** decoder, from development to deployment. Deployed 11 configurations across backends (**AIMET, TorchAO**) and targets (RTX 4060, **Qualcomm QCS8550 Hexagon HTP**), delivering one at 36% lower on-device latency retaining 99% embedding cosine. (First-author preprint with failure-mode catalog; [repo](#))
- **Synthetic data generation and orchestration:**
A controllable **compositional data synthesis** framework for **domain transfer** and **bias evaluation**, targeting **driver monitoring** and **defect detection**. Provides **cross-environment** event and behaviour **transfer** via decoupling-and-recombination to tackle **data scarcity**. Proposal under review by **Google TPU Research Cloud**. In parallel, as an engineering PoC, built a **multi-agent LLM orchestration** service on **AWS** (ECS, S3, Docker), using ChromaDB retrieval and HDBSCAN clustering to cut redundant API calls.

Doctoral Researcher | IDLab, Ghent University-imec, Sep 2019 - Nov 2025

Funded by the Flanders AI Research Program (Flemish government initiative on data-efficient and reliable AI).

- **Real-world stress test of SOTA audio models on 5000h of unlabeled streams:**
With industry partner **AsaSense**, **stress-tested** SOTA **audio models** under real deployment across **multiple sites** (2 cities, 5 sensors) and **time** (multiple years) periods with uncurated urban audio. Architected a context-aware, **model-agnostic** acoustic surveillance system combining **cloud-side event tagging** and **on-device anomaly detection**, **reducing** cloud **workload** by **90%**. Identified **lab-to-field gaps** including **domain shift**, **privacy** exposure, and **open-set** events; summarized findings into a technical report for our industry partner.
- **Single-model audio-visual representation for multi-task edge deployment:**
Built a **general-purpose, self-supervised** representation learning framework supporting **multi-task** downstream **applications**, allowing **efficient deployment** on **resource-constrained edge devices**. Recovered **false negatives** in **contrastive learning** for more information-rich representations, yielding **~10% gain** over baselines. Tackled the **gap** between **current benchmark assumptions** and **real-world** audio-visual streaming **data**. (Published in *Frontiers in Robotics and AI*)
- **On-device opt-in audio privacy without downstream model retraining:**
Implemented the **first opt-in** framework in the **voice privacy** domain and extended it to visual modality, deployable as a **plug-and-play** module with **real-time inference** demonstrated on **Jetson TX1** (34 ms per 1s clip). Two **CycleGAN-VC2**-based models engage in **adversarial learning** to **reduce** privacy **leakage** to near **zero**. Challenged the **overlooked gap** between **opt-out** assumption in **benchmarks** and **opt-in** principle of EU **General Data Protection Regulation**. (Published in *IEEE Pervasive Computing and Applied Intelligence*)

- **Model selection for general-purpose transfer without sharing training data:**

Proposed a **task-agnostic, source-free, unsupervised transferability assessment** for efficient **domain adaptation** prior to **model fine-tuning**. Forms unbiased reference embeddings with an **ensemble** of **randomly** initialised neural **networks**, then ranks pretrained models by embedding similarity, achieving **0.95/1** against fully supervised ranking without source data and target label. Bridged the **overlooked gap** between **EU Data Act's** model-sharing mandate and **industry** concerns over **data sovereignty** and **IP**. (Published in *Sensors*)

Teaching Assistant (Applied Machine Learning) | Ghent University, *Sep 2021 - Feb 2023*

- **Collaborated with Sony DepthSensing Solutions:**

Co-designed a [RGB-D hand gesture recognition project](#) with **Sony DepthSensing Solutions**, from depth sensor data acquisition to PoC delivery.

Pre-PhD Research | NCKU and Academia Sinica, *Sep 2014 - Jun 2016, Sep 2017 - Jul 2019*

C++, Tensorflow, OpenCV, MATLAB

- **Turned research into an award-winning prototype, pitched to industry and academic juries:**

Spun off the research direction into a **startup pitch**. Engineered the core **prototype (video event analysis)** and presented business plans to industry and academic juries; won **1st** at InnoServe Awards and **Top 3** at APICTA.

- **Facial video analysis for driver monitoring with national automotive R&D partner:**

Built a **driver monitoring system** fusing **facial video descriptors** with **ECG** signals, in collaboration with **ARTC** (Automotive Research and Testing Center, Taiwan). (Published in *IEEE TITS*, **60+** citations).

- **Vision research at National Cheng Kung University (NCKU) and Academia Sinica:**

Conducted research on **trajectory clustering** and **superpixel**. (Published in *IEEE TIP*; *IEEE ICIP*)

Selected [Publications](#)

13 peer-reviewed (**150+ citations**): **4 first-author journals** (2 Top 10%, 1 Top 25%), **4 first-author conferences** (2 oral), 3 co-author journals (2 Top 10%, 1 Top 25%), and 1 co-author workshop.

[Source-free model transferability assessment for smart surveillance via randomly initialized networks](#)

Sensors, 2025 (journal, **first-author**)

[Embedding-based pair generation for contrastive representation learning in audio-visual surveillance data](#)

Frontiers in Robotics and AI, 2025 (journal, **first-author**)

[Privacy-preserving visual analysis: Training video obfuscation models without sensitive labels](#)

Applied Intelligence, 2024 (journal)

[Selective manipulation of disentangled representations for privacy-aware facial image processing](#)

4th Workshop on Machine Learning for CyberSecurity at ECML PKDD, 2022 (workshop)

[An opt-in framework for privacy protection in audio-based applications](#)

IEEE Pervasive Computing, 2022 (journal, **first-author**)

[Driver monitoring using sparse representation with part-based temporal face descriptors](#)

IEEE Transactions on Intelligent Transportation Systems (TITS), 2020 (journal, **60+ citations**)

[Clustering trajectories in heterogeneous representations for video event detection](#)

IEEE International Conference on Image Processing (ICIP), 2018 (conference, **first-author, oral**)

[USEAQ: Ultra-fast superpixel extraction via adaptive sampling from quantized regions](#)

IEEE Transactions on Image Processing (TIP), 2018 (journal)

[Event based surveillance video synopsis using trajectory kinematics descriptors](#)

IAPR International Conference on Machine Vision Applications (MVA), 2017 (conference, **first-author**)

[Spatiotemporal coherence-based annotation placement for surveillance videos](#)

IEEE Transactions on Circuits and Systems for Video Technology (TCSVT), 2016 (journal, **first-author**)

[Video gender recognition using temporal coherent face descriptor](#)

IEEE/ACIS SNPD, 2015 (conference, **first author**)

[Trajectory kinematics descriptor for trajectory clustering in surveillance videos](#)

International Symposium on Circuits and Systems (ISCAS), 2015 (conference, **first-author, oral**)

Education

Ph.D. in Computer Science Engineering, Ghent University, Belgium (*Sep 2019 - Nov 2025*)

Master of Science in Computer and Communication Engineering, NCKU, Taiwan (*Sep 2014 - Jun 2016*)

Bachelor of Science in Electrical Engineering, NCKU, Taiwan (*Sep 2010 - Jun 2014*)