Yang-Che (Jeff) Tseng

Taipei, Taiwan

Email : jeff@yangchetseng.com
GitHub · LinkedIn · Google Scholar

ML infrastructure engineer with research in 3D computer vision (NeurIPS, ECCV, ICRA, WACV) and hands-on experience running large-scale GPU experiments on HPC clusters, shipping LLM applications, and building data pipelines in industry.

Education

- **University of British Columbia** — Average 92.8%
  M.Sc. in Computer Science – Advisor: Kwang Moo Yi — *Sep 2024 – Nov 2026*
- **National Taiwan University** — GPA 3.92/4.0
  B.Sc. in Computer Science and Information Engineering — *2019 – 2023*

Experience

- **University of British Columbia** — Vancouver, Canada
  M.Sc. Thesis Research, 3D Computer Vision — *Sep 2024 – Nov 2026*
  Research Assistant, Computer Vision Lab — *2023 – 2024*
  - **Evaluation protocol:** found that single-image NVS benchmarks score misaligned predictions; proposed a fix (matching → PnP → co-visibility mask) and re-benchmarked **11 baselines** on 4 datasets (1.7k test views).
  - **GPU fleet at scale:** ran **110k+ Slurm jobs (≈15.8k GPU-hours)** on 3 clusters (A5000/3090, V100, H100 MIG), peak **127 concurrent GPUs**; templated per-scene jobs, packaged baselines as **Apptainer** images.
  - **Agentic HPC ops:** built a **Claude Code** setup (48 skills, 6 subagents, safety hooks, post-mortem rules) that submits, monitors and fixes Slurm jobs over SSH; parallel autonomous sessions in worktrees opened **281 PRs (227 merged)**.
  - **Throughput:** cut CogNVS eval from ∼3.5 h to ∼21 min per scene by batching across V100s.
  - **Methods:** MapAnything-GS (MapAnything SfM init → **3D Gaussian Splatting**); as RA, dynamic-scene 3DGS (optical-flow static/dynamic split, video-diffusion refiner) and PointNeRF++ experiments.
- **MyProGuide (TravoAI)** — Remote
  Software Engineer (Part-time) — *Jul 2024 – Mar 2025*
  - **Serverless email pipeline:** built an **AWS SES → Lambda → S3** pipeline that parses cancellation emails from 12 OTAs (Klook, KKday, Viator, . . .) and updates orders via API; **in production 1+ year** at up to 35k emails/month.
  - Wrote multilingual per-OTA parsers with a **Jest** suite on real emails; recovered 3,200 of 3,600 historical emails with Ops.
  - **Flutter app MVP:** built a tour-guide app from scratch to MVP via spec-first issues and reviewed PRs (BLoC, OAuth, CI).
  - **LLM tooling:** prototyped a **text-to-SQL** agent (LangChain, Gemini) that creates PostgreSQL test accounts.
- **GoGoX** — Taipei, Taiwan (remote from Vancouver)
  Software Engineer, Data Scientist — *Jun 2023 – Jun 2024*
  Data Scientist Intern — *Sep 2022 – Jun 2023*
  - **LLM QA bot:** built a RAG chatbot (**OpenAI, LangChain, Pinecone**) on **GCP**; cut latency 2 s with **Redis Stack**.
  - **Driver stopping model:** extracted the **GraphHopper** routing engine from a deprecated Java dispatch system into a **Docker Compose** service with a Python interface; ran counterfactual analysis on driver stop probability.
  - **Labeling automation:** ML model on GPS/order data to auto-label pickup times, with **autoencoder** anomaly filtering.
  - **GPS map matching:** **HMM** matching to OpenStreetMap (Python, shapely) to build a travel-time database for ETA.
- **ASUS Intelligent Cloud Services (AICS)** — Taipei, Taiwan
  Software Engineering Intern — *Jul 2022 – Sep 2022*
  - Extended Selenium IDE with **Electron / Node.js**, making 19 of 23 previously unrecordable test cases automatable.

Publications

- **WACV 2027** – *Aligning Misaligned Benchmarks for Single-View Novel View Synthesis.* **Y.-C. Tseng**, E. Dong, J.-g. Kwak, G. Sharma, S. Kheradmand, S. Garbin, H. Isack, K. M. Yi.
- **NeurIPS 2024** – *3D Gaussian Splatting as Markov Chain Monte Carlo.* S. Kheradmand, D. Rebain, G. Sharma, W. Sun, **Y.-C. Tseng**, H. Isack, A. Kar, A. Tagliasacchi, K. M. Yi.
- **ECCV 2024** – *PointNeRF++: A Multi-Scale, Point-Based Neural Radiance Field.* W. Sun, E. Trulls, **Y.-C. Tseng**, S. Sambandam, G. Sharma, A. Tagliasacchi, K. M. Yi.
- **ICRA 2023** – *Orbeez-SLAM: A Real-time Monocular Visual SLAM with ORB Features and NeRF-realized Mapping.* C.-M. Chung, **Y.-C. Tseng**, Y.-C. Hsu, X.-Q. Shi, Y.-H. Hua, J.-F. Yeh, W.-C. Chen, Y.-T. Chen, W. H. Hsu.

Skills & Leadership

- **Languages:** Python, C/C++, CUDA, JavaScript/TypeScript (Node.js), Dart
- **ML & Vision:** PyTorch, 3D reconstruction (SfM, Gaussian Splatting, NeRF, SLAM), LLM apps (RAG, agents)
- **Infrastructure:** Slurm, Apptainer, HPC, AWS (Lambda, SES, S3), Google Cloud, Cloudflare Workers, GitHub Actions
- **Spoken:** Mandarin (native), English (fluent)
- **TA**, NTU Deep Learning for CV (2022), DSA (2021); **Lead, Wired Network Team**, NTU CSIE NASA (2022–23).
- **Open source:** added per-user filtering to **nvtop**, a GPU monitoring tool (C).