Fanyi Pu

+65 8398-1454 | FPU001@e.ntu.edu.sg | pufanyi.com | 58tv6skAAAAJ | pufanyi

Education

Nanyang Technological University — Singapore
Bachelor of Science in Data Science and Artificial Intelligence — *2022 – 2026*

- Expected Honours (Highest Distinction), Current CGPA: 4.63 / 5.00.
- Research Interests: **Spatial Intelligence** and **Unified Multimodal Models**. 1300+ Google Scholar citations.
- Core courses: Reinforcement Learning, Deep Learning, Data Structure and Algorithms, Probability and Statistics, Computational Economics, Stochastic Process, Numerical Analysis, Cryptography, Electricity & Magnetism.

University of California, Berkeley — Berkeley, CA
Summer Session — *Jun 2024 – Aug 2024*

- Studied *Computer Security* and *Game Theory*, GPA: 4.00 / 4.00

Publications & Research Projects

SenseNova-SI: Scaling Spatial Intelligence with Multimodal Foundation Models
Preprint, Co-first author — *Sep 2025 – Present*

- A foundation model designed to scale spatial intelligence, reached state-of-the-art on key spatial benchmarks.
- Conducted rigorous ablation studies to establish that spatial intelligence is primarily driven by data scaling rather than reasoning heuristics (CoT), guiding the project's strategic focus on constructing the SenseNova-SI-8M dataset.
- Executed the full-stack training pipeline for Qwen-based variants with LMMs-Engine on 100+ GPUs.

LMMs-Engine: A Simple, Unified Multimodal Models Training Engine
Research Project, Core Developer — *Aug 2025 – Present*

- A lean and flexible training framework designed for both rapid research prototyping and large-scale production.
- Optimized the Bagel training pipeline by integrating FSDP2 and Liger Kernel for efficient distributed training.

LMMs-Eval: Reality Check on the Evaluation of Large Multimodal Models
NAACL (Findings), Co-first author — *Jan 2024 – Present*

- Multimodal evaluation frameworks with 3.3K GitHub stars.
- Built the pipeline code; Developed the low-cost automatic generation pipeline for the Multi-modal LiveBench, which leverages continuously updating news and online forums to evaluate models' generalization capabilities in the wild.

Video-MMMU: Evaluating Knowledge Acquisition from Multi-Discipline Professional Videos
Preprint, Third Author — *Sep 2024 – Feb 2025*

- Built an evaluation set with 300 expert-level videos and 900 human-annotated questions across six disciplines, cited by Gemini 3 Pro (Google) and GPT-5 (OpenAI).
- Proposed knowledge gain metric quantifies improvement in performance after watching the video lectures; Validated the integrity and accuracy of expert-level data in economics and medicine, ensuring benchmark quality.

VLoRP: Memory-Efficient LLM Training by Various-Grained Low-Rank Projection of Gradients
Preprint, Fourth author — *Aug 2024 – Jan 2025*

- Implemented the training framework with DeepSpeed integration and reproduced baselines (LoRA, GaLore, MeZO).

Multi-Modal In-Context Instruction Tuning (Otter, MIMIC-IT)
IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Co-first author — *Jun 2023 – Dec 2023*

- Early experiment on a vision-language-agent (VLA) model with 3.3k stars on GitHub.
- Developed a pipeline to generate a multimodal instruction tuning dataset using pure language models, which generated 2.8M data for training; Created a multimodal evaluation framework, which later evolved into the LMMs-Eval.

Experiences

Research Intern (Spatial Intelligence) — Sep 2025 – Present
SenseTime Research — *Singapore*

- Conducted research on **spatial intelligence**, focusing on the optimization of unified multimodal model.
- Delivered **SenseNova-SI**, establishing a new state-of-the-art model for spatial understanding.

Core Member — Jan 2024 – Present
LMMs-Lab — *Singapore*

- Core member of a non-profit initiative dedicated to democratizing Large Multimodal Models (LMMs); instrumental in the lab's 0-to-1 establishment, from initial ideation and naming to its current operations.
- Spearheaded the development of high-impact open-source projects: the evaluation framework **LMMs-Eval**, the efficient inference library **LMMs-Engine**, and the video understanding benchmark **VideoMMMU**.

Core Contributor Jan 2025 – Jul 2025
Synvo AI *Singapore*

- Architected and implemented the **Synvo File System**, the company's inaugural product, serving as the storage backbone for unstructured multimodal data.
- Empowered the core contextualization engine by structuring data for efficient retrieval, enabling the development of AI systems that can remember, learn, and adapt.

Research Intern (Multimodal Models) Jan 2025 – May 2025
MMLab@NTU *Singapore*

- Supervised by Prof. Liu Ziwei, focused on **multimodal language models** and **unified multimodal models**.

Competitions

The International Collegiate Programming Contest (ICPC) 2021 – 2024
Nanyang Technological University / Dalian University of Technology *Singapore / China*

- Ranked 22th in 2024 ICPC Asia Pacific Championship.
- Ranked 13th in 2023 ICPC Asia Jakarta Regional Contest.
- Ranked 6th in 2022 ICPC Asia Manila Regional Contest.
- Gold medal in 2021 ICPC Asia Kunming Regional contest.
- Silver medal in 2021 ICPC Asia Nanjing Regional contest.

Simon Marais Mathematics Competition Dec 2022
Nanyang Technological University *Singapore*

- Best-in-University Prize in Nanyang Technological University.
- Solved problems in number theory, game theory, and calculus.

Teaching & Activities

Tutorial Lecturer Jan 2026 – Present
NTU SC1008 C & C++ PROGRAMMING *Singapore*

- Instructed a cohort of 50 beginners in core C/C++ syntax and methodologies.

Nanyang Programming Contest Jan 2025 – Present
Nanyang Technological University *Singapore*

- Organized 5 competitions and tutorials to non-ICPC students.
- Aimed at helping students struggling with algorithm problems in interviews.

NTU Students' Computing and Data Science Club Nov 2024 – Present
Nanyang Technological University *Singapore*

- Managed a telegram group with 130+ members.
- Answer questions about AI/ML and share learning resources related to deep learning research.

Teaching Assistant Aug 2024 – Dec 2024
NTU SC1003 Introduction of Computational Thinking *Singapore*

- Instructed a cohort of 20 beginners in core Python syntax and data manipulation methodologies.

Hobbies

Physics, Chess (lichess rating: 2084), Calligraphy, Table tennis