

FRANK WU

(516) 817-2165 | frankwu2@andrew.cmu.edu | [LinkedIn](#)

Education

Carnegie Mellon University Pittsburgh, PA
Master of Science in Machine Learning Dec 2026

- Cumulative GPA: 3.9 / 4.0; Courses Taken: Advanced Machine Learning, Probability & Statistics

New York University New York, NY
Bachelor of Arts in Joint Computer Science & Mathematics Jun 2024

- Cumulative GPA: 3.8 / 4.0; Honors: Cum Laude; International Collegiate Programming Contest (ICPC) Top 21
- Campus Activities: Grader (Data Structures), Co-organizer (NYU ProgTeam), Startup Co-Founder (Bamboost)
- Courses Taken: Graduate Machine Learning, Heuristic Problem Solving, Linear Algebra, Honors Probability, Analysis, Numerical Analysis, Algebra

Publications

Frank Wu, Mengye Ren. *Local reinforcement learning with action-conditioned root mean squared Q-functions*. International Conference on Learning Representations (ICLR), 2026. <https://arxiv.org/abs/2510.06649>.

Research Interests

Exploring emerging paradigms for training neural networks, with an emphasis on reinforcement learning, continual learning, and generalization. Current interests include mechanisms that enable iterative computation and refinement, and how such dynamics can support reasoning.

Research Experience

Courant Institute, NYU New York, NY
Action-conditioned Root-mean-Squared Q-Functions (ARQ) Aug 2024 – Sep 2025

- Proposed ARQ, unifying Forward-Forward learning with TD updates for fully local RL; accepted to **ICLR 2026**
- Introduced action-conditioned, vector-valued value functions, substantially increasing capacity in local architectures
- Achieved 50% average return improvement via action conditioning at model input
- Evaluated on MinAtar and DeepMind Control Suite; outperformed prior local-RL methods and matched backprop-based baselines

Column Forward-Forward (ColumnFF) May 2023 – May 2024

- Proposed ColumnFF, achieving 99.55% on MNIST and 76.32% on CIFAR-10 without backpropagation; submitted to NeurIPS 2024
- Replaced layer normalization with instance normalization, improving CIFAR-10 accuracy by 3%
- Designed supervised label embedding, improving accuracy by an additional 1%

Work Experience

Alibaba Group Hangzhou, China
Machine Learning Engineer, Future Life Lab Apr 2025 – Aug 2025

- Built *Shop with Me* feature for Pailitao, a visual search engine in Taobao serving 50M+ DAU
- Curated 200K labeled samples based on Qwen3-235B; performed SFT distillation on private models achieving 98% safety, 95% readability, 87% fidelity
- Integrated fashion-specific domain knowledge into SOTA multimodal models, reaching 90% precision and 80% accuracy on in-house benchmarks
- Built MVP using FastAPI and JavaScript with fashion style analysis and VLM-based filtering modules

Machine Learning Engineer, Rhino Intelligent Manufacturing Jul 2024 – Mar 2025

- Built a visual action recognition foundation model for apparel manufacturing based on InternVideo2, achieving 90% accuracy on internal benchmarks
- Developed *ActionTAL*, a temporal action segmentation algorithm based on AdaTAD, reaching 75.3 mAP
- Deployed ML modules for Shengong customization platform with 80% accuracy in 30-shot settings
- Led curation of a 900K-video dataset for domain-specific pretraining, including public and in-house data

Skills

Programming Languages: Python, C++, Java, MATLAB, SQL, C, JavaScript, C#, R, Swift, LaTeX
Frameworks & Tools: PyTorch, JAX, TensorFlow, NumPy, Pandas, Scikit-learn, Linux, Slurm