

Yucheng (Steven) Chen

stevenyc.chen@mail.utoronto.ca | schen.me | github.com/perctrx | Toronto, ON, Canada

Education

University of Toronto

Computer Science, Faculty of Arts & Science (University College)

Toronto, Canada

Sept. 2026 – 2030 (expected)

Elton Academy

Ontario Secondary School Diploma

Ontario, Canada

2025

Publications & Preprints

- Yucheng Chen. **Synthetic Benchmarks Overstate Forward-Forward Scaling: Real-Data Limits of Layer-Local Training**. *arXiv preprint*, Jun. 2026.
- Tianhao Fu, Austin Wang, Charles Chen, Roby Aldave-Garza, Yucheng Chen. **SegWithU: Uncertainty as Perturbation Energy for Single-Forward-Pass Risk-Aware Medical Image Segmentation**. *arXiv preprint; under review at UNSURE 2026 (MICCAI workshop)*, Apr. 2026.
- Tianhao Fu, Yucheng Chen. **MIP Candy: A Modular PyTorch Framework for Medical Image Processing**. *arXiv preprint*, Feb. 2026.
- Yucheng Chen. **ATLAS-CNN: A Novel Hybrid Deep Learning Architecture for Stock Market Prediction**. *SSRN preprint*, Jan. 2025.

Experience

Research Intern

Apr. 2025 – Present

Project Neura / UTMIST

- Research on medical image analysis, focused on uncertainty and reliability
- Co-first-author on MIP Candy; co-author on SegWithU

AI / Full-Stack Engineer Intern

Feb. 2025 – Present

Kabuda Group, Toronto

- Part-time. Kabuda runs a multi-service platform for newcomers and international students in the GTA.
- Built conversational AI features and agent workflows for client-facing web and mobile products, and the REST backends they run on

Founder

2025 – Present

Amplimit

- Home for my Forward-Forward scaling work and for Ampdraft

Projects

Forward-Forward Scaling Audit (DTG-FF) | Python, PyTorch

2024 – 2026

- Built DTG-FF (dynamic temperature goodness, decoupled normalization, multi-layer fusion) and evaluated it across nine real-data benchmarks, ahead of every FF variant it was compared against
- Produced the first FF-family baseline at ImageNet-100 224x224, and used the instrument to show that architecture-matched backprop controls widen their lead as class count grows
- Showed that synthetic teacher-student K-sweeps reverse sign against real images, and that FF's $O(1)$ -in-depth activation memory does not beat gradient accumulation on commodity hardware

MIP Candy | Python, PyTorch

2025 – Present

- Modular framework for medical image processing: a full training, inference, and evaluation pipeline from a single `build_network` method
- LayerT, a deferred configuration mechanism that swaps convolution, normalization, and activation modules at runtime without subclassing
- k-fold CV, automatic ROI dataset inspection, deep supervision, EMA, and training-state recovery. Apache-2.0 on PyPI; 294 stars, 44 forks

SegWithU | Python, PyTorch

2026

- Post-hoc uncertainty estimation that models uncertainty as perturbation energy, letting a frozen segmentation model flag its own failures in a single forward pass
- No retraining, no ensembling, and no repeated sampling required

IV Surface Explorer | Python, Dash, Plotly, MCP

2026

- IV surface, smile, term structure, skew, GEX, open interest walls, max pain, and expected move over normalized full option chains
- Exposed as an MCP server so a language model can query the pipeline directly

Ampdraft | Next.js, Prisma, PostgreSQL, Typst/WASM

2026

- Version control for prose: snapshots, parallel variants, and word-level diff, deliberately avoiding Git vocabulary
- Client-side PDF export via Typst compiled to WebAssembly

Technical Skills

Languages: Python, Java, C++, Rust, TypeScript **ML/DL:** PyTorch, computer vision, medical image segmentation, uncertainty estimation **Web & Data:** Next.js, React, NestJS, Prisma, PostgreSQL, Redis **Infrastructure:** Docker, Nginx, Linux, Git