YIFENG WANG

Email | *Scholar*

EDUCATION

Carnegie Mellon University — *Pittsburgh, USA*
Ph.D. in Electrical and Computer Engineering — *Sep. 2024—May 2029 (expected)*
Advisor: Prof. Sarah Cen
M.S. in Electrical and Computer Engineering — *Sep. 2024—May 2026*
(Fritsch Family Fellowship 2024–2025)

The Hong Kong Polytechnic University — *Sep. 2019—Jun. 2023*
Bachelor of Science in Computer Science, Minor in Applied Mathematics
(Dean's Honor List for 2021-2022)

PUBLICATIONS

1. **Linear Dynamical Modeling of Language Models**
Yifeng Wang, Sarah Cen
Under review at *ICLR'27*.

2. **AutoAL: Automated Active Learning with Differentiable Query Strategy Search** (*Paper* | *Code*)
Yifeng Wang, Xueying Zhan, Siyu Huang.
ICML'25.

3. **Rapid Image Labeling via Neuro-Symbolic Learning** (*Paper* | *Code*)
Yifeng Wang*, Zhi Tu*, Yiwen Xiang, Shiyuan Zhou, Xiyuan Chen, Bingxuan Li and Tianyi Zhang
KDD'23 (Oral Presentation).

4. **SonicSieve: Bringing Directional Speech Extraction to Smartphones Using Acoustic Microstructures** (*Paper*)
Kuang Yuan*, **Yifeng Wang***, Xiyuxing Zhang*, Chengyi Shen, Swarun Kumar, Justin Chan
CHI'26.

5. **ADAPT: Adaptive Profiling Transformers for Efficient Alzheimer's Disease Diagnosis** (*Paper*)
Yifeng Wang, Ke Chen, Haohan Wang
MICCAI'26.

6. **One Shot for All: Quick and Accurate Data Aggregation for LPWANs** (*Paper*)
Ningning Hou, Xianjin Xia, **Yifeng Wang**, Yuanqing Zheng
INFOCOM'23 (Extended version: IEEE/ACM Transactions on Networking).

7. **MoLoRa: Intelligent Mobile Antenna System for Enhanced LoRa Reception in Urban Environments** (*Paper*)
Ningning Hou, **Yifeng Wang**, Xianjin Xia, Shiming Yu, Yuanqing Zheng, Gu Tao.
SenSys'25.

8. **HEPHA: A Mixed-Initiative Image Labeling Tool for Specialized Domains** (*Paper* | *Code*)
Shiyuan Zhou, Bingxuan Li, Xiyuan Chen, Zhi Tu, **Yifeng Wang**, Yiwen Xiang and Tianyi Zhang.
IUI'25.

9. **A Unified Evaluation-Instructed Framework for Query-Dependent Prompt Optimization** (*Paper*)
Ke Chen, **Yifeng Wang**, Hassan Almosapeeh, Haohan Wang
Under review at *JMLR*.

10. **DS-ViT: Dual-Stream Vision Transformer for Cross-Task Distillation in Alzheimer's Early Diagnosis** (*Paper*)
Ke Chen, **Yifeng Wang**, Yufei Zhou, Haohan Wang.

11. **From Tissue Plane to Organ World: A Benchmark Dataset for Multimodal Biomedical Image Registration using Deep Co-Attention Networks** (*Paper*)
Yifeng Wang, Weipeng Li, Thomas Pearce, Haohan Wang

SELECTED RESEARCH EXPERIENCE

Linear Dynamical Modeling of Language Models
Sep. 2025—Sep. 2026
Carnegie Mellon University | Advisor: Prof. Sarah Cen

• **Contribution:** We proposed an LDS-based framework for characterizing LLM hidden-state evolution through time-lagged covariance dynamics, revealing interpretable modes of generation stability.

• **Approach:** We developed low-dimensional trajectory summaries that capture temporal structure in both synthetic systems and real LLM generations.

• **Achievement:** We demonstrated applications to mathematical reasoning verification, hallucination detection, and jailbreak susceptibility, outperforming several LLM-as-a-judge baselines at negligible additional inference cost.

Directional Speech Extraction with Acoustic Microstructure (CHI'26)
Oct. 2024—Apr. 2025
Carnegie Mellon University | Advisor: Prof. Justin Chan, Co-Advisor: Prof. Swarun Kumar

• **Contribution:** We designed **the first directional speech extraction system on the smartphone** with a user-friendly acoustic microstructure.

• **Approach:** We optimized a bio-inspired acoustic microstructure design for spatial awareness. We built an off-the-shelf in-line microphone product which can be easily adapted to multiple kinds of smartphones. We built a real-time angle-based speech extraction network with a specially designed angle encoder.

• **Achievement:** We demonstrated that our system achieves a performance of **5dB sisnr improvement** in unseen environments, which is better than **conventional 5-microphone arrays**.

Differentiable Active Learning (ICML'25)
Mar. 2024—Oct. 2024
Carnegie Mellon University | Advisor: Prof. Siyu Huang

• **Contribution:** We designed the first differentiable AL strategy search method, which can select the best samples.

• **Approach:** Our design consists of two neural networks, which are iteratively co-optimized in a differentiable bi-level optimization framework using the labeled data, learning how well a set of candidate AL algorithms perform on that task.

• **Achievement:** Our approach consistently achieves superior accuracy compared to all candidate AL algorithms and other selective AL approaches, with at most **3%** improvement in **all AL iterations**.

Rapid Labeling of Images for Specialized Classification Tasks (KDD'23, IUI'25)
May. 2022—Sep. 2024
Purdue University | Advisor: Prof. Tianyi Zhang

• **Contribution:** We designed an image labeling system to substantially reduce the number of human-annotated images required for specialized image classification tasks (e.g. glaucoma detection, bird species classification).

• **Approach:** We leverage existing object detection algorithms to produce bounding boxes and text labels of important features within the image (e.g. bird mouth, feet, wings), and use this as input to an inductive programming system that is able to infer the image class (e.g. orange crowned Warbler). The system also provides a human-interpretable explanation for the output (e.g. black eyes, white feet, orange wings)

• **Achievement:** We demonstrated that using a modest dataset of **39 images**, our approach achieves a performance of **88%** across four image labeling tasks, which is **15% higher** than conventional CNN-based approaches.

Detecting Alzheimer's Disease using 3D MRI Brain Scans (MICCAI'26)

Mar. 2023—Nov. 2023
University of Illinois Urbana-Champaign | Advisor: Prof. Haohan Wang

- **Contribution:** We designed a Transformer-based neural network model to detect Alzheimer's from a patient's 3D MRI brain scans.
- **Approach:** Our design consists of four different encoders and a fusion-attention mechanism to extract and combine visual information across different slices and sections of the 3D volume.
- **Achievement:** Our approach achieves an accuracy of **92%**, which is **8%** higher than SOTA models designed for the task, while using 47% less parameters.

Optimizing Wireless Node Placement Using Robotic Arms (SenSys'25) *Jun. 2022—Oct. 2023*
Hong Kong PolyU | Advisor: Prof. Yuanqing Zheng

- **Contribution:** Designed a robotic system to search for the optimal wireless receiver position and orientation which would maximize SNR within a region of an unseen environment with unknown multipath characteristics.
- **Approach:** We created a Bayesian optimization model that would take three signal calibration points sampled in the environment to predict the wireless channel within a 1 m sphere.
- **Achievement:** Our system achieves an SNR **improvement of 7 dB** over prior works in the domain.

AWARDS & RECOGNITION

Fritsch Family Fellowship 2024-2025
Carnegie Institute of Technology Dean's Fellow 2024-2025
Dean's Honor List for 2021-2022 (Top 5% student)
Hong Kong Blockchain Developer Hackathon Awardee 2021
Final Year Project Award (Scholarship: $4000)

SERVICE

Reviewer for NeurIPS 2023, 2025, 2026, AAAI 2024, 2025, MICCAI 2025, 2026, CVPR 2026

PROFESSIONAL SKILLS

Programming: Python, C/C++, Java, SQL, Shell, MATLAB
ML Frameworks: PyTorch, TensorFlow, scikit-learn
Research: LLMs, representation learning, reinforcement learning, optimization, statistical modeling