YIXIONG CHEN

Baltimore U.S.A.
+1 (410) 310 1390 ◇ ychen646@jh.edu

EDUCATION

Johns Hopkins University (JHU) Baltimore, United States
Ph.D. in Computer Science (Overall GPA: 4.00 / 4.00) *Sep. 2023 - May 2028 (Expected)*

- Working on VL foundation models and generative models under the supervision of Alan Yuille.
- Interested in multi-modal large language models, diffusion models, and self-supervised learning.

Chinese University of Hong Kong (CUHK) Shenzhen, China
Ph.D. in Data Science (Overall GPA: 3.90/4.00) *Sep. 2021 - Jul. 2023*

- Excellent Paper on Science and Technology of Shenzhen (2023)
- Excellent Paper on Artificial Intelligence of Shenzhen (2022)

Fudan University Shanghai, China
B.Sc. in Data Science (Major GPA: 3.42/4.00, Overall GPA: 3.31/4.00) *Sep. 2016 - Jun. 2021*

- Highest Impact Award of IEEE Computer Society Biometrics Workshop (2024)
- Second Class Scholarship for Outstanding Graduates (2021)
- Huawei Cloud Scholarship (2018)

PROFESSIONAL EXPERIENCE

Google Research Mountain View, United States
Student Researcher, Health AI Team *Jun. 2024 - present*

· Curated a chest X-ray images-report pair longitudinal dataset from MIMIC-CXR, Chest Imagenome, and MS-CXR-T datasets to contain multiple images with disease progression description.
· Optimized vision encoders for contrastive captioner (CoCa) models by designing a local attention mechanism for multi-image fusion, achieving 4% improvement on temporal classification.
· Wrote a technical report with extensive experiments. To be submitted to **MICCAI 2025**.

Johns Hopkins University Baltimore, United States
Research Intern, CCVL (Alan Yuille's Lab) *Jul. 2022 - Mar. 2023*

· Analyzed the training process of MLPs, CNNs, and ViTs, on 8 datasets, and discovered a fundamental learning preference called *layer convergence bias*. The results were published at **ICLR 2023**.
· Explored the radiomic features of liver tumors in CT slides and proposed a morphological algorithm for 3D tumor synthesis, which can pass the Visual Turing Test conducted by experienced clinical doctors. The results were published at **NeurIPS workshop 2022** and **CVPR 2023**.

Shenzhen Research Institute of Big Data (SRIBD) Shenzhen, China
Research Assistant, Leader of Medical Group *Jul. 2020 - Jul. 2023*

· Forged a label selection/correction algorithm inspired by the KNN representation space for effective learning with noisy labels. The results were published at **ICASSP 2024**.
· Devised powerful layer-wise fine-tuning algorithms for pre-trained medical models based on Meta-learning. The results were published at **MICCAI 2023**.
· Developed and perfected pre-training algorithms for ultrasound videos with contrastive learning. The results were published at **MICCAI 2021** and **TMI 2022**.
· Designed novel models for liver fibrosis staging based on lesion localization and multi-modal training.

Points Technology Shanghai, China
Machine Learning Intern *Mar. 2020 - Jun. 2020*

- Designed and implemented the SecureBoost algorithm with Numpy, which is an improvement of XGBoost under the horizontal federated learning standard.
- Optimized the SecureBoost to reduce its time complexity from $O(n^2)$ to $O(nlogn)$.
- Participated in the formulation of the standard *"Promotion Committee of China Communication Standardization Association: Federated learning technology tool for data circulation"*.

RESEARCH PROJECTS

Vision Language Understanding for Longitudinal Chest X-Ray Google, Summer 2024

- Analyzed chest x-ray datasets, proposed a data pipeline to break down reports into symptom description and progress judgment for vision-language alignment and progression prediction respectively.
- Devised a **regional cross-attention** mechanism to enhance the vision encoder's local comparison performance to recognize the nuanced difference of five conditions between longitudinal image pairs.
- Explored a **multi-stage fine-tuning** strategy to synchronize the pre-trained and randomly initialized parameters. The model outperformed previous SOTA by 8.6% accuracy on temporal classification.

Medical Image Segmentation Evaluation Tools JHU, Spring 2024

- Proposed a novel **text-embedding-conditioned** label (mask) quality evaluation tool, Quality Sentinel, for medical image segmentation performance evaluation without reference.
- Built a large-scale dataset that contains CT scans paired with pseudo labels of varying quality and their corresponding ground truth DSC. The data size is **4M image-mask pairs**.
- Conducted extensive experiments to show Quality Sentinel's performance in diagnosing dataset quality, human-in-the-loop annotation, and semi-supervised segmentation.

Human-centric Text-to-Image Generation JHU, Autumn 2023

- Constructed a large-scale human dataset containing three subsets: human-in-the-scene images, close-up of face images, and close-up of hand images, totally 1,000,000 images.
- Proposed a novel **Mix of Expert (MoE)** mechanism to fuse two low-rank experts (LoRA) for enhancing the generation of hands and faces, respectively.
- Evaluated the model on two proposed novel benchmarks based on COCO and DiffusionDB human prompts, and reached **SOTA** among all open-sourced text-to-image diffusion models.

Text-to-Image Generation Benchmarking CUHK, Spring 2023

- Proposed a novel **hierarchical Chain of Thought** to prompt MiniGPT-4 for explainable evaluation of images generated by text-to-image diffusion models.
- Built a large-scale **human preference dataset** on images from a wide range of sources comprising 800k human preference choices on 430k pairs of images.
- Used COCO Caption and DrawBench as prompt sets to benchmark 15 models' capabilities to generate images based on texts *w.r.t.* their performance of fidelity, alignment, and aesthetics.

Liver Tumor Synthesis JHU, Autumn 2022

- Proposed a **morphological liver tumor synthesis** pipeline based on location selection, texture/shape generation, and image warping, to generate 49,000 examples with annotated synthetic tumors.
- Combined clinical knowledge to the algorithm for a random hyper-parameter tuning scheme, enabling the synthetic tumors to be diverse and to cheat doctors.
- Conducted nnUNet and Swin UNETR tumor segmentation with hybrid synthetic/realistic tumors to achieve **SOTA on the LiTS** dataset.

Basic Layer Properties of DNNs CUHK and JHU, Summer 2022

- Defined a measurement for transferability based on optimization path, and found the **layer-wise transferability** trend from ImageNet to 12 downstream image classification tasks.
- Defined a measurement for **layer-wise convergence rate**, analyzed the learning speeds for different layers in a DNN, and found shallower layers tend to learn faster than deeper layers.

Transfer Learning for Medical Image Analysis CUHK, 2020 - 2022

- Proposed a **meta-learning**-based fine-tuning method to automatically determine transfer strengths for different layers of a DNN, reaching 2%+ performance improvement.
- Implemented a **meta-contrastive learning** framework for ultrasound videos, leveraging the semantic clustering of 3000 videos to reach more than 10% performance gain than ImageNet pre-training.

Person Re-identification with Changing Clothes Fudan University, Spring 2019

- Proposed a person re-id benchmark including 23,000+ pedestrians, each with 4 different clothes.
- Collected, organized, and annotated a video dataset for pedestrians using GTA-V virtual engine.
- Tested 10 existing SOTA person re-id algorithm frameworks on our benchmark.

PUBLICATIONS

Click to view Google Scholar

1. Zhu J, **Chen Y**, Ding M, Luo P, Wang L, Wang J. MoLE: Human-centric Text-to-image Diffusion via Mixture of Low-rank Experts. **NeurIPS 2024.**
2. Li W, Qu C, Chen X, Bassi P, Shi Y, Lai Y, Yu Q, Xue H, **Chen Y**, et al. AbdomenAtlas: A large-scale, detailed-annotated, & multi-center dataset for efficient transfer learning and open algorithmic benchmarking. **Medical Image Analysis 2024.**
3. **Chen Y**, Zhou Z, Yuille A. Quality Sentinel: Estimating Label Quality and Errors in Medical Segmentation Datasets. **Arxiv 2024 preprint.**
4. Jiang H, **Chen Y**, Ding C, Liu L, Han X, Zhang X. Leveraging Noisy Labels of Nearest Neighbors for Label Correction and Sample Selection. **ICASSP 2024.**
5. **Chen Y**, Li J, Jiang H, et al. MetaLR: Meta-tuning of Learning Rates for Transfer Learning in Medical Imaging. **MICCAI 2023.**
6. Wu X, Hao Y, Sun K, **Chen Y**, Zhu F, Zhao R, Li H. Human Preference Score v2: A Solid Benchmark for Evaluating Human Preferences of Text-to-Image Synthesis. **Arxiv 2023 preprint.**
7. **Chen Y**, Liu L, Ding C. X-IQE: eXplainable Image Quality Evaluation for Text-to-Image Generation with Visual Large Language Models. **Arxiv 2023 preprint.**
8. Hu Q, **Chen Y**, Yuille A, Zhou Z. Label-Free Liver Tumor Segmentation. **CVPR 2023.**
9. **Chen Y**, Yuille A, Zhou Z. Which Layer is Learning Faster? A Systematic Exploration of Layer-wise Convergence Rate for Deep Neural Networks. **ICLR 2023.**
10. **Chen Y**, Li J, Liu L, Ding C. Rethinking Two Consensus of the Transferability in Deep Learning. **Arxiv 2023 preprint.**
11. **Chen Y**, Zhang C, Liu L, Ding C. Generating and Weighting Semantically Consistent Sample Pairs for Ultrasound Contrastive Learning. **IEEE Transactions on Medical Imaging.**
12. Hu Q, Xiao J, **Chen Y**, Sun S, Chen JN, Yuille A, Zhou Z. Synthetic Tumors Make AI Segment Tumors Better. **NeurIPS 2022 Workshop.**

13. Zhang C, **Chen Y**, Liu L, et al. HiCo: Hierarchical Contrastive Learning for Ultrasound Video Model Pretraining. **ACCV 2022.**
14. **Chen Y**, Zhang C, Liu L, et al. Uscl: Pretraining deep ultrasound image diagnosis model through video contrastive representation learning. **MICCAI 2021 (oral).**
15. Wan F, Wu Y, Qian X, **Chen Y**, Fu Y. When person re-identification meets changing clothes. **CVPR 2020 Workshop.**

PRESENTATIONS

Meta-Learning-Rate for Medical Transfer Learning (Poster) Vancouver, Canada
The 26th Int. Conf. on Medical Image Computing and Computer Assisted Intervention *Oct. 2023*

Layer-wise Convergence Rate of DNNs (Poster) Kigali, Rwanda
The 11th International Conference on Learning Representations *May. 2023*

Meta-Reweighting for Contrastive Learning (Poster) Shenzhen, China
The 2nd Doctoral & Postdoctoral Academic Forum at SRIBD *Aug. 2022*

Ultrasound Contrastive Learning (Oral) Online
The 24th Int. Conf. on Medical Image Computing and Computer Assisted Intervention *Sept. 2021*

ACDEMIC ACTIVITIES

Reviewer: NeurIPS, ICLR, MICCAI, TASE	2024
Reviewer: MICCAI, ICML, TASE	2023
Teaching Assistant: AI in Medical Imaging	2022 Spring
Teaching Assistant: Statistical Inference	2021 Fall

RELEVANT COURSES

Core Courses	**Other Courses**
Machine Learning	Cloud Computing
Deep Learning and Neural Networks	Data Visualization
Advanced Convex Optimization	Probability and Statistics
Algorithms for Numerical Computation	Linear Algebra
Computer Vision	Advanced Mathematics

SKILLS

Programming	Python, Bash, Matlab, R, JavaScript
Tech Stack	Linux, Git, Slurm, PyTorch, JAX, Numpy, Pandas, OpenCV, LaTeX
Language	English (TOEFL 107), Chinese (mother tongue)