

Senhao (Simon) Cheng

senahoc@umich.edu | [Google Scholar](#) | [Personal Website](#)

Research Interests: Multimodal Reasoning, Vision-Language Models, Reinforcement Learning for VLMs, Cross-modal Retrieval & Generation

Education

University of Michigan, Ann Arbor M.S. in Electrical and Computer Engineering	Ann Arbor, USA 2024 – 2026 (expected)
Zhejiang University B.Eng. in Automation	Hangzhou, China 2020 – 2024

Preprints & Submissions

Decompose, Look, and Reason: Reinforced Latent Reasoning for VLMs Mengdan Zhu*, <u>Senhao Cheng*</u> , Liang Zhao (*Equal Contribution) <i>Under Review at ACL Rolling Review, 2026</i>	Ann Arbor, USA Jan 2026 – Present
--	--------------------------------------

- Led the design of the visual grounder module (two cross-attention layers + FFN, 32 learnable latent queries) and the two-pass forward mechanism that injects premise-conditioned visual latents into the VLM’s reasoning chain at each step.
- Formulated and implemented **SGLP** (Spherical Gaussian Latent Policy), which extends GRPO to continuous latent space via ℓ_2 -projection onto the unit hypersphere, resolving the magnitude collapse observed with standard Gaussian exploration.
- Built the full three-stage training pipeline end-to-end (InfoNCE pretraining on VQAv2 \rightarrow LoRA SFT with rank = 64 on Vision-R1-cold \rightarrow Dr. GRPO + SGLP) using PyTorch and DeepSpeed on 2 \times H200 GPUs.
- SOTA on four benchmarks with Qwen3-VL-8B, surpassing GPT-4o (~200B): V* Bench 83.8%, MathVista 67.5%, MMMU-Pro 56.1%, MMStar 65.2%.

Cross-modal RAG: Sub-dimensional Text-to-Image Retrieval-Augmented Generation

 Ann Arbor, USA

Mengdan Zhu*, <u>Senhao Cheng*</u> , Guangji Bai, Yifei Zhang, Liang Zhao <i>Under Review at ACL Rolling Review, 2026</i> arXiv:2505.21956	Mar 2025 – Present
---	--------------------

- Designed and trained the sub-dimensional dense retriever: a lightweight cross-attention adaptor on top of frozen CLIP ViT-L/14 that produces subquery-conditioned vision embeddings at 0.01 \times CLIP’s GPU memory cost.
- Implemented the Pareto-optimal retrieval algorithm guaranteeing all non-dominated images are returned under multi-objective retrieval.
- Ran all retrieval experiments on MS-COCO and Flickr30K; integrated GPT-image-1 and Gemini-2.0-flash as generation backbones for the full RAG pipeline.
- MS-COCO R@1 81.82% (prev. best 59.10%), Flickr30K R@1 97.50%; consistent generation gains on WikiArt, CUB, and ImageNet-LT.

ChemSafetyBench: Benchmarking LLM Safety on Chemistry Domain

 Beijing, China

Haochen Zhao*, Xiangru Tang*, ..., <u>Senhao Cheng</u> , ..., Mark Gerstein <i>Preprint</i> arXiv:2411.16736	Apr 2024 – Nov 2024
---	---------------------

- Contributed to hazardous substance data collection and quality control; ran adversarial evaluation pipelines (name-hacking, CoT attacks) across 10+ LLMs.
- Benchmark covers 30K+ QA pairs over 1,700 substances; revealed that GPT-4o scores near chance on chemical property identification.

Publications

A Breast Cancer Detection Model Based on Modified ConvNeXt v2 <u>Senhao Cheng</u> , Esther Sun, Wangzi Qian, Yang Han <i>AIBDF 2023, ACM</i> DOI: 10.1145/3660395.3660493	Irvine, USA Jul 2023 – Sep 2023
--	------------------------------------

- First author. Replaced global average pooling with Generalized-Mean Pooling and swapped AdamW for AdaBelief in ConvNeXt v2; ran all experiments and wrote the paper. pF1 improved by 0.031–0.043 over

ResNet50, GoogLeNet, and EfficientNet-B2 on mammography classification.

Industry Experience

MindRank Ltd.

Hangzhou, China

AI & Data Analysis Intern

Sep 2023 – Apr 2024

- Wrote Python ETL pipelines to clean, deduplicate, and load 5TB+ biomedical records (diseases, drugs, patents, clinical trials) into the PharmKG knowledge graph.
- Trained Random Forest and Gradient Boosting classifiers on structured biomedical features to flag drug-target correlations for internal drug discovery workflows.

Technical Skills

Frameworks & Models: PyTorch, HuggingFace Transformers, DeepSpeed, CLIP, Diffusion Models, Qwen-VL

Methods: Reinforcement Learning (GRPO/PPO), LoRA Finetuning, Retrieval-Augmented Generation, Contrastive Learning, Multi-objective Optimization

Tools & Infra: Python, C, MATLAB, SQL, Git, Linux, AWS, Weights & Biases, vLLM, TensorFlow