

Pengcheng (Ryan) Xu

+1(217)550-1337 | pengcheng.xu@uci.edu | <https://github.com/explcre> | <https://linkedin.com/in/pengcheng-xu-ryan>
Google Scholar: <https://scholar.google.com/citations?user=LXzbdZEAAAAJ&hl=en>
Personal Website: <https://pengcheng-xu.com> | ChatGPT version of me website: <https://pengcheng-xu.com/mychat/>

EDUCATION

University of California, Irvine

Sep 2024 – 2029 (expected)

Ph.D. in Computer Science. GPA: 4.0/4.0

Research Interest: Large Language Model, Reinforcement Learning, On-Policy Distillation, Knowledge Distillation, Conditional & Controllable Generative Models, Multi-modal Language Model, Bioinformatics, AI for Science, Large Language Model for Science, DNA Language Model, Single Cell, Spatial Transcriptomics Foundation Model.

Courses: AI in Biology and Medicine(A), Innovation and Entrepreneurship(A), Light and Geometry in Computer Vision(A+), Deep Generative Models(A).

University of Illinois Urbana-Champaign

Aug 2022 - Dec 2023

Master's Degree in Computer Engineering, GPA: 3.8/4.0. Research Scientist at [AI@UIUC](#).

Core Courses: Distributed Systems, Communication Network(A+), Computer Vision(A), Advanced Distributed Systems, Transfer Learning(A), Engineering Entrepreneurship(A+), Machine Learning(A), Applied Parallel Programming.

University of Michigan - Shanghai Jiao Tong University Joint Institute

Sep 2018 - Aug 2022

Bachelor's Degree in Electrical and Computer Engineering. Minor in Data Science. Outstanding Graduate (School Level).

Core Courses: Data structures and Algorithms, Intro to Computer Organization (A, Teaching Assistant), Computational Methods for Statistics and Data Science (A+), Intro to Operating Systems, Intro to Data Science (A), Undergraduate Research (A+).

PUBLICATIONS

(Manuscript in preparation) Pengcheng Xu, Junhao Liu, Peter Tan Dat Le, Haoyi Cai, Bruhati Aarushi Kuchi, Arnav Abhijit Dhariya, Chaoyang Wang, Hokin Deng, Jing Zhang. "DNAThinker: A Reasoning-Grounded Multimodal Foundation Model for Cell-Type-Conditioned DNA Generation, Editing, and Pair Prediction with Mask-Diffusion Reinforcement Learning".

(Under review at *Science Advances*) Pengcheng Xu*, Junhao Liu*, Yi Dai*, Kainoa Andrew Nishida, Dongbo Sun, Yutong Lei, Yaqi Hu, Chaoyang Wang, Jing Zhang. OmegaGenome: Toward Better Sub-Million Scale Expert Models from Large Genomic Language Models via Knowledge Distillation. (* equal contribution.)

(Published in the KDD 2026 proceedings) Simon Dongbo Sun, Junhao Liu, Pengcheng Xu, Yaqi Hu, Martin Jinze Zhang, Jing Zhang. [MUGO: Differentiable Combinatorial Optimization for Causal Variant Discovery in the Non-coding Genome](#), pp. 12042–12053.

(Accepted to ICML 2026) Maijunxian Wang, ..., Pengcheng Xu, Hokin Deng. [A Very Big Video Reasoning Suite](#). [Paper]

(Manuscript in preparation) Wenbo Zhang*, Pengcheng Xu* (co-first), et al. "Divergence-Aware On-Policy Distillation for Multi-Turn Long-Horizon Autonomous LLM Agents" (* equal contribution).

(Submitted to AAAI 2027) Pengcheng Xu, et al., Jing Zhang. "Prediction Instability, Not Confidence, Routes Test-Time Compute in Recursive Reasoners and Language Models".

(Submitted to AAAI 2027) Pengcheng Xu, et al., Jing Zhang. "From Wasserstein Distances to Wasserstein Barycenters: Prediction-Space Ensembling in Set-Valued Knowledge Distillation".

(Submitted to the August 2026 ACL Rolling Review cycle) Pengcheng Xu, et al. [TexGemma-D1: An LLM Training Pipeline with a Verifiable Symbolic Reward for Scientific Diagram Generation](#). Pandita AI / diagram.ai.

(Manuscript in preparation) Pengcheng Xu. "Stability Is Not Value: Budget-Coupled Compute Allocation for Vision–Language–Action Models".

(Manuscript in preparation) Pengcheng Xu. "Selection Without Signal: Heritability, Allocation, and a Price List for Self-Evolving Agents".

(Manuscript in preparation) Pengcheng Xu. "VERA: Verification-Rewarded Test-Time Adaptation for Agentic Scientific Data Analysis".

(Under review) Yutong Lei, Yaqi Hu, Chaoyang Wang, Pengcheng Xu, Junhao Liu, Henry Cheng, Jing Zhang. SPATREG: A Hierarchical Framework for Histology Image Registration in Spatial Transcriptomics.

(Manuscript in preparation) Xi Li, Ziheng Duan, Diptanshu Sikdar, Pengcheng Xu, Jing Zhang. BRIDGE: Spatially Aware Bimodal Representation Integration with Hierarchical Graphs for Spatial Transcriptomics Resolution Enhancement.

(ACL BioNLP Workshop 2025) Pengcheng Xu, Sicheng Shen, Jieli Zhou, Hongyi Xin. [Team XSZ at BioLaySumm2025: Section-Wise Summarization, Retrieval-Augmented LLM, and Reinforcement Learning Fine-Tuning for Lay Summaries](#), pp. 275–280.

(ACM BCB 2024, SIGBIO Paper Award at BCB '24) Ziqi Rong, Jinpu Cai, Jiahao Qiu, Pengcheng Xu, Lana X. Garmire, Qiuyu Lian, Hongyi Xin. [L2 Normalization and Geodesic Distance for Enhanced Information Preservation in Visualizing High-dimensional Single-cell Sequencing Data](#), pp. 1–11.

(Accepted to KDD-AIDSH Workshop 2024, Oral) Pengcheng Xu, Tianfan Fu, Wenhao Gao, Jimeng Sun. [REINVENT-Transformer: Molecular De Novo Design through Transformer-based Reinforcement Learning](#), 2024.

(ACL BioNLP Workshop 2024) Jieli Zhou*, Cheng Ye*, Pengcheng Xu, Hongyi Xin. [Team YXZ at BioLaySumm: Adapting Large Language Models for Biomedical Lay Summarization](#), pp. 818–825. (* equal contribution by the first two authors.)

(arXiv:2306.13866v2) Pengcheng Xu*, Jinpu Cai*, Yulin Gao, Ziqi Rong. [MIRACLE: Multi-task Learning based Interpretable Regulation of Autoimmune Diseases through Common Latent Epigenetics](#), 2023. (* equal contribution by the first two authors, as defined in

the public v2 source.)

(Poster presented at CSHL Single Cell Analyses 2023) Jinpu Cai, Ziqi Rong, Luting Zhou, Xinzhu Jiang, Pengcheng Xu, Rui Gao, Yu Zhao, Bing He, Jianhua Yao, Qiuyu Lian, Hongyi Xin. "Cell-ontology-aware marker identification and stratification with TAMER".

(Unpublished technical report) Pengcheng Xu, Kaiyang Chen, Yuanrui Zhang, Indranil Gupta. [Pipe-Dejavu: Hardware-aware Latency Predictable, Differentiable Search for Faster Config and Convergence of Distributed ML Pipeline Parallelism](#), 2023.

Zhikai Yang, Pengcheng Xu, Dekun Yang, Yufeng Chen, YanCong Ma. [Vascular Intervention Training System Based on Electromagnetic Tracking Technology](#), ICVRV 2020, pp. 269–270.

RESEARCH EXPERIENCE

DNAThinker: A Reasoning-Grounded Multimodal Foundation Model for Cell-Type-Conditioned

Feb 2026 - Present

DNA Generation, Editing, and Pair Prediction

Lead researcher & primary author. Manuscript in preparation.

- ▷ Built DNAThinker, a unified regulatory-DNA foundation model coupling a Qwen3.5-0.8B + LoRA reasoning LLM with an NTv3-100M-post + MDLM DNA generator through an SD3/FLUX-style cross-attention bridge and delta-residual conditioning path over a 1.5M-row fine-tuned regulatory-DNA corpus.
- ▷ Developed tool-grounded reasoning prompts with FIMO motif scans and per-cell expression context, so rationales cite verifiable transcription factors and JASPAR motifs across T1 enhancer generation, T2 pair prediction, and T3 enhancer editing.
- ▷ Applied strict-holdout grounded-RL post-training whose reward pairs a TF-grounding term with an explicit hallucination penalty, and measured it as a head-to-head against the supervised (SFT) checkpoint: on a scaled 7-cell holdout (n=700) under the refined v2 detector, RL lifts TF-grounding from the SFT baseline of 0.702/0.714/0.799 to 0.919/0.956/0.920 on T1/T2/T3 (e.g. $\sim 0.80 \rightarrow 0.92$ on T3 enhancer editing) while cutting hallucinated TFs by 74%/67%/67% – isolating the hallucination-reward RL step, not SFT, as the source of the grounding gain.
- ▷ Introduced vector-level TF-knockdown intervention on the 17,415-d cell-expression slot plus knockdown-aware policy-gradient training, achieving 63% conditional motif depletion on the Mic T3 test slice and 56% depletion / 71% depleting rows across the 7-cell stratified aggregate.
- ▷ Developed MDLM-RL, a policy-gradient framework for masked-diffusion DNA generators with GRPO, DAPO, GSPO, and SVGSPO variants; audited reward, algorithm, capacity, and learning-rate regimes, reporting T1 single-shot FID 8.78 (4.4x better than LEONINE), a 14x edit-tight T3 trunk-only motif lift, and clear saturated-trunk/no-RL regimes.

Divergence-Aware On-Policy Distillation for Multi-Turn Long-Horizon Autonomous LLM Agents

2026 - Present

Co-first author (listed second; equal contribution with Wenbo Zhang). Manuscript in preparation. Improves on TCOD / vanilla OPD.

- ▷ Co-developing a divergence-aware on-policy distillation (OPD) method for multi-turn, long-horizon autonomous LLM agents that improves on vanilla OPD and the temporal-curriculum baseline TCOD, using a lightweight runtime student-teacher disagreement signal – computed from log-probabilities already produced by the OPD forward pass, at negligible added cost – to adaptively govern the per-step distillation objective so the student stays within the teacher’s reliable support, stopping the trajectory-level KL collapse that destabilizes OPD over long horizons.
- ▷ Reproduced the full TCOD multi-turn-agent distillation stack on Trinity-RFT (verl 0.7 / vLLM / Ray) over ALFWorld and WebShop: verified the distillation teacher is the GiGPO Qwen2.5-7B-RL ALFWorld checkpoint and reproduced the paper’s teacher row to the decimal (85.71% seen / 76.12% unseen success), then SFT/cold-started and distilled Qwen2.5-3B students under vanilla OPD and the TCOD-B2F (back-to-front) / F2B (front-to-back) curricula, reproducing the headline OOD-generalization result (F2B 79.1% unseen vs. paper 79.19, +13 pts over OPD) under a faithful paper-HP harness across seen/unseen/hard splits, and surfacing that a more stable reproduced OPD baseline (seen 82.1 vs. paper 65.7) compresses TCOD’s reported margin.
- ▷ Engineered the multi-GPU SLURM training/eval infrastructure (free-GPU pinning, FSDP+SDPA path without flash-attn, cross-node SSD remapping, teacher-logprob OOM fixes) and an inference-only ALFWorld eval workflow with verified zero-missing game-file data integrity across all splits; root-caused and fixed a trajectory memory-append bug in the curriculum training code whose history-buffer handling mis-aligned the on-policy distillation window used by the teacher-prefix and B2F/F2B schedules.
- ▷ Extending the benchmark suite beyond ALFWorld/WebShop to tool-integrated-reasoning (TIR) and search-agent settings – adding the Search-R1 retrieval-agent benchmark and the SOD (step-wise on-policy distillation) TIR baselines on AIME 2024/2025, GPQA-Diamond, and LiveCodeBench; faithfully reproduced the released SOD-1.7B to within ~ 3 pts of the paper on AIME (average@32 48.1/36.8 vs. 50.8/41.7 under the model’s native sampling) to validate the head-to-head comparison harness.

Prediction Instability, Not Confidence, Routes Test-Time Compute in Recursive Reasoners and Language Models

2026 - Present

Lead researcher. Submitted to AAAI 2027.

- ▷ Established that *prediction instability* – how much a prediction moves under a label-preserving perturbation – not the model’s *confidence*, is the right label-free, training-free signal for routing test-time compute across recursive reasoners (Tiny Recursion Model, HRM) and LLMs (Qwen2.5/3.5, Phi-3.5, Llama-3.1): confident-but-wrong predictions are stable fixed points, so confidence mis-routes – entropy-based halting collapses to depth 1.8 / accuracy 0.199 where instability holds 0.859, and rescue-prediction AUROC is 0.88 (instability) vs. 0.48 (confidence) on GSM8k.
- ▷ Introduced a *free* latent-velocity signal that gates the depth axis – certifying on 3/3 tasks whether extra recursion pays (Sudoku alive, ARC/Maze converged) – and showed across 17 test-time per-step interventions that nothing beats taking the deepest step on a converged recursion; composing extended depth with augmentation-consistency gating (ACG) lifts Sudoku-Extreme exact-grid

accuracy from 0.82 to 0.992 at depth-96 (+0.12 to +0.18 over the trained model, zero retraining).

- ▷ Built VR-MoR, a training-free Mixture-of-Recursion that routes per-token recursion depth by velocity – 3.8x less compute at near-full accuracy, dominating the uniform-depth frontier by up to +0.108 and matching the model’s *trained* ACT halt head with no training, labels, or auxiliary network.
- ▷ Extended the same signal to a generalist→specialist model-capacity cascade (Qwen2.5-3B→7B; always-specialist accuracy at 73% compute, rescue-AUROC 0.83 vs. confidence 0.50) that transfers down the price ladder (3B→7B 0.74, 0.5B→7B 0.62), a unified cross-axis allocator that Pareto-dominates any single axis at every budget, and a routed self-refinement lever that exceeds refine-all at 38% of its cost.
- ▷ Built a realizable online controller (one calibrated threshold, no labels or reward model) that reaches the oracle accuracy–compute frontier (oracle accuracy at ~14x vs. 66x uniform on a mixed Sudoku+Maze workload; 75% of self-consistency’s gain at 44% of its samples on GSM8k); independently reproduced HRM-Text-1B’s published benchmark table (5 matches), found and ablated a PrefixLM-masking pitfall inflating naive HF evaluation by +0.36, and beat the 2025 SOTA confidence method DeepConf head-to-head (routing AUROC 0.666 vs. chance).

From Wasserstein Distances to Wasserstein Barycenters: Prediction-Space Ensembling in Set-Valued Knowledge Distillation

2026 - Present

Lead researcher. Submitted to AAAI 2027.

- ▷ Developed prediction-space ensembling for set-valued knowledge distillation from a frozen 108.7M-parameter single-cell perturbation teacher into independently seeded 3.21M-parameter students; a five-student pool improved the primary Pearson- Δ metric by +0.0179 against an exactly output-size-matched control, improving all 70/70 held-out perturbations ($t=13.6$).
- ▷ Derived metric-geometry results explaining why affine-invariant, first-moment evaluation limits the direct headroom of reweighting and higher-moment objectives, and when pooled predictions and a W_2 barycenter score equivalently; tested ten objective variants to convergence with same-seed pairing and found seven ties, three losses, and no positive arm in the fixed benchmark.
- ▷ Measured 26% of the converged student’s residual as seed-averageable functional variance ($R^2=0.999$); matched ten independent members with five runs times two stored snapshots at half the training compute, while explicitly separating prediction pooling from cross-basin weight averaging.
- ▷ Validated the method across domains on CIFAR-100: a six-student prediction pool gained +2.42 top-1 and improved calibration, and distilling the pool into one student recovered +1.20 top-1 across 6/6 seeds at single-model inference; also identified seed-derived evaluation and emitted-output-count confounds worth up to 0.0177 and 0.34, respectively.

Stability Is Not Value: Budget-Coupled Compute Allocation for Vision–Language–Action Models

2026 - Present

Lead researcher. Manuscript in preparation; preregistered study.

- ▷ Proposed Memory-Guided Exit (MGE), a history- and budget-conditioned gate that allocates per-episode VLA compute by predicting the marginal value of more depth from stall recurrence, compute debt, remaining budget, and proprioception instead of action stability or confidence.
- ▷ Built simulator-forked counterfactual depth rollouts and a latency-coupled LIBERO/LIBERO-Plus protocol in which the environment advances during inference, plus swappable memory, oracle, and contextual-bandit modules with 42 tests; benchmark results are pending.

Selection Without Signal: Heritability, Allocation, and a Price List for Self-Evolving Agents

2026 - Present

Lead researcher. Manuscript in preparation. Public [AReAL research branch](#).

- ▷ Reconstructed 198 evaluations from 225 public tree snapshots spanning 24 MGM candidates and found no detectable between-candidate held-out signal in that released run: raw means beat the grand mean in 1.5% of 4,000 split-halves, while James–Stein shrinkage won 98.5%, despite search expanding high-mean nodes ($r=0.29$).
- ▷ Built mutation-tested AReAL instrumentation, held-out evaluation, outcome-calibrated critics, and matched controls; in a controlled Qwen2.5-1.5B GRPO run, all six trained checkpoints fell below the held-out base model and training reward mis-ranked capability, motivating evidence-based promotion gates.

VERA: Verification-Rewarded Test-Time Adaptation for Agentic Scientific Data Analysis

2026 - Present

Lead researcher. Manuscript in preparation.

- ▷ Reproduced GeneBench-Pro and BioMysteryBench harnesses and audited 73 human-solvable tasks: doubling the step budget from 25 to 50 flipped 0/20 graded failures, while model-authored checks constrained 0/17 audited cases, falsifying the initial budget and self-verification hypotheses.
- ▷ Implemented a typed, bounded state graph with executable harness invariants, provenance, forced-answer gating, trajectory voting, and fixed-model evaluation; maintained 127 research, 70 BioMysteryBench, and 82 GeneBench-Pro tests, with all 25 recorded mutations caught.

A Very Big Video Reasoning Suite: Public Task Generators and Dataset Construction

2025 - 2026

Contributing researcher. Accepted to ICML 2026; paper PDF: <https://arxiv.org/abs/2602.20159>.

- ▷ Implemented eight public Python prototype data-generator repositories for [gear-chain direction](#), [arrow maze](#), [scale-balance tilt](#), [dot quantity comparison](#), [layer-depth ordering](#), [water-level reasoning](#), [binary counter](#), and [odd-one-out](#).
- ▷ Built prototype generation and validation logic mapping symbolic task states into controllable video prompts/assets and answer labels.
- ▷ Assisted with Wan video-generation-model fine-tuning experiments for dataset generation.

OmegaGenome: Toward Better Sub-Million Scale Expert Models from Large Genomic Language Models via Knowledge Distillation.

Mar 2025 - Present

Lead researcher & primary author (I wrote most of the code + paper).

- ▷ Designed and implemented OmegaGenome, a two-stage knowledge distillation framework evaluated on 18 classification tasks and a separate base-resolution multi-track regression benchmark. Across the current manuscript's five classification teachers under three-seed evaluation, distilled experts are up to 28,750x smaller, retain up to 98% of teacher task performance, and retain 95% on average.
- ▷ Fine-tuned all four teacher architectures using full-parameter training for Enformer, DNABERT-2, and Caduceus, and LoRA (rank=48) for NT-2.5B, reducing trainable parameters to <1% of the 2.5B backbone; distilled teachers into dilated CNN students (BPNetClassifier, 2K-3.6M params) via a composite objective combining KL divergence, MSE feature alignment, and cross-entropy.
- ▷ Benchmarked distilled students against teacher models and size-matched baselines, demonstrating 31-1,671x inference speedup, 7.5-513x GPU memory reduction, and consistent 6.8-10.8% MCC improvement over scratch-trained BPNet across epigenetic mark prediction, enhancer/promoter classification, and splice site detection.
- ▷ Implemented and compared four KD algorithms (OmegaGenome, DKD, DIST, Logit Standardization) with systematic grid search over 45+ configurations; quantified representation transfer via cosine similarity, t-SNE/UMAP, and logit-level R2 (distilled student 0.779 vs. scratch baseline 0.677), confirming faithful feature inheritance from teacher to student.
- ▷ For the Science Advances revision, added a billion-scale *generative* teacher – Carbon-3B, a Llama-architecture genomic LM – and distilled it into a 0.12M-parameter BPNet student (~25,000x smaller) across the 18 NT benchmark tasks; a single-config distillation (CE/KL/MSE 0.5/0.5/0.2) already retains 93.4% of the teacher's mean MCC, exceeds the size-matched from-scratch baseline overall, and beats the teacher outright on the two noisiest histone marks, with per-task hyperparameter tuning (a 1,440-config grid) and 3-seed confirmation in progress.
- ▷ Reproduced the base-resolution sequence-to-function foundation model NTv3-650M-post (InstaDeep; single-nucleotide tokenization, up to 1 Mb context) on its 34-track human benchmark to paper-grade fidelity – per-track Pearson 0.606 averaged over 3 seeds (std ≤ 0.003), matching or exceeding the published model on ATAC-seq, histone ChIP-seq, and PRO-cap – establishing a strong billion-scale, quantitative-readout teacher for the revision's base-resolution experiments.
- ▷ Extended OmegaGenome from scalar classification to *base-resolution, per-base-pair multi-track regression* (histone-mark, chromatin-accessibility, PRO-cap, eCLIP, and RNA-seq coverage signals across cell types), designing a DilatedTrackNet student whose ~32 kb receptive field covers the full input window, and benchmarking SOTA regression-KD objectives (Poisson-multinomial, DIST, channel-wise, logit-standardized MSE) for compressing NTv3 into compact (~1–8M) deployable per-bp students.

MUGO: Differentiable Combinatorial Optimization for Causal Variant Discovery in the Non-coding Genome

2025 - 2026

Contributing researcher (ran benchmark experiments). Published in the KDD 2026 proceedings, pp. 12042–12053.

- ▷ Contributed to MUGO, a framework that recasts non-coding causal-variant discovery as differentiable combinatorial optimization: a sparse variant-selection mask is relaxed with a Gumbel-Softmax distribution and a straight-through estimator, keeping DNA inputs discrete for a frozen sequence-to-signal oracle (Borzoï, 524kb context) while gradient ascent searches the combinatorial variant space, with a 10-head Monte-Carlo consensus for variance reduction.
- ▷ Ran benchmark experiments across 5 molecular modalities (RNA-seq, CAGE, ATAC-, ChIP-, DNase-seq) and 7 tissues against ISM (greedy/exhaustive), saliency maps, random search, CADD and FunSeq2; MUGO delivered the largest predicted expression gain (e.g., 174.8 vs. 58.2 saliency on blood RNA-seq) at near-constant ~2 min/locus, while a separate memory analysis showed the ensemble fit within 24GB, versus the intractable $O\left(\binom{N}{k}\right)$ subset search of exhaustive ISM.
- ▷ Evaluated prioritized variants against GTEx v8 cis-eQTLs, where MUGO achieved 2.50x enrichment over random versus 1.94/0.97/1.20 for the reported baselines; separately evaluated GWAS Catalog variants, confirmed cross-backbone robustness (Borzoï vs. Enformer, $r=0.56$), and ran an LD-resolution stress test separating $PIP>0.9$ causal variants from high-LD ($r^2>0.8$) proxies.

Team YXZ at BioLaySumm: Adapting Large Language Models for Biomedical Lay Summarization

Apr 2024 - June 2024

- ▷ Evaluated prompt-only OpenBioLLM-Llama3-70B, whose test-set LENS score was 81.2052 versus 48.0963 for the BART baseline (68.8% relative improvement); the team ranked first on the shared task's readability measure.
- ▷ Fine-tuned Llama3-8B using LoRA; on the reported validation comparison, LENS increased from 80.1033 to 80.7012, AlignScore from 0.6316 to 0.70628, and SummaC from 0.5282 to 0.5816.
- ▷ Implemented semantic-similarity K-shot prompting; at $K=1$, the reported automated factuality metrics improved by 15.0% relative for AlignScore and 7.9% relative for SummaC.

REINVENT-Transformer: Molecular De Novo Design through Transformer-based Reinforcement Learning

May 2023 - May 2024

Advisors: Tianfan Fu (then tenure-track Assistant Professor at Rensselaer Polytechnic Institute; now tenure-track Associate Professor at Nanjing University) and Jimeng Sun (Professor of Computer Science, UIUC).

- ▷ Implemented a transformer-based autoregressive molecular generator with oracle-feedback reinforcement learning; in the public paper's AUC Top-10 comparison, the method achieved the highest aggregate score (12.197) and the best or tied-best result on 8 of 20 tasks.
- ▷ Applied Oracle-feedback reinforcement learning on the downstream tasks to reach higher performance than pretrained model.
- ▷ Carried out ablation study and investigation into loss curve and conditional probability over the next token as a function of previously chosen ones according to the model.

Hardware-aware Latency Predictable, Differentiable Search for Faster Config and Convergence of Distributed ML Pipeline Parallelism

Feb 2023 - May 2023

Advisor: Indranil Gupta, Professor of CS UIUC | Advanced Distributed Systems | Researcher.

- ▷ Implemented a predictive model that considers communication cost, model computational cost, and hardware information to predict latency and resources of parallel configurations, saving time on pre-profiling before searching the parallel configuration.
- ▷ Proposed a differentiable parallel configuration search space inspired by DARTS, can potentially reach optimal configuration faster than the original dynamic programming.
- ▷ Employed parallel random initialization using sampling algorithms like Bayesian Optimization for faster train loss convergence.

Multi-modal target detection with zero-shot depth estimation, Multi-modal NAS

May 2022 - Sep 2022

Shanghai AI Lab | Intelligent Photoelectric Department | Multi-modal Cognitive Computing Algorithm Intern.

- ▷ Literature review on nature/science papers on multi-modal learning, and conference papers on Neural Architecture Search, multi-modal learning, image and sound feature fusion.
- ▷ The model inputs image and 2 audio channels, and outputs the behavior category and distance. Improve the estimation of distance using zero-shot monocular depth estimation like MiDaS added to the model structure.

MIRACLE: Multi-task Learning based Interpretable Regulation of Autoimmune Diseases through Common Latent Epigenetics | Research Assistant

Sep 2021 - Dec 2023

Advisor: Hongyi Xin, Associate Professor of UM-SJTU Joint Institute, Shanghai Jiao Tong University.

- ▷ Explored an adaptable, interpretable neural network for learning common DNA-methylation patterns across approximately 480K CpG sites from hundreds of samples.
- ▷ Designed a site-gene-pathway constrained neural architecture whose learned weights can prioritize candidate sites, genes, and pathways for follow-up.
- ▷ Implemented a variational autoencoder with a pathway-level latent embedding shared across disease datasets for multi-task learning.
- ▷ Optimized the reported pretrain-fine-tune scheme, increasing the six-dataset mean accuracy from 84.10 to 87.20 (+3.10 percentage points, approximately 3.7% relative), and contributed to writing the public 2023 arXiv v2 preprint.

Augmented reality simulation of cardiovascular interventional surgery | Research Assistant

Mar 2020 - May 2021

Advisor: Lixu Gu, Professor of Biomedical Engineering, Shanghai Jiao Tong University.

- ▷ Developed the framework of an augmented reality surgery training assistant system for medical student and surgery.
- ▷ Predicted the operation trajectory using LSTM and used KD-Tree to calculate the distance for operation safety warning.
- ▷ Displayed vascular model in AR with OpenGL and designed the UI interface to support translation.
- ▷ Used the aruco library in OpenCV to coordinate positioning of the QR code.
- ▷ Published Vascular Intervention Training System Based on Electromagnetic Tracking Technology on ICVRV as second author.

Normal fundus image generation based on Generative Adversarial Network | Research Assistant

Apr 2019 - Oct 2019

Advisor: Lisheng Wang, Professor at SEIEE, Shanghai Jiao Tong University.

- ▷ Utilized GAN to design an image generation system to generate normal fundus image for comparisons to identify lesion areas.
- ▷ Used datasets including MESSIDOR-2 for data preprocessing like normalization, contrast enhancement, de-noising, clipping.
- ▷ Implemented a U-Net-based segmentation network to identify and remove fundus blood vessels as part of de-noising process.
- ▷ Built a CycleGAN framework based on segmentation to generate the normal fundus image to identify the lesion area.

PROFESSIONAL EXPERIENCE

Pandita AI Inc. (diagram.ai) | Large Language Model Research Scientist Intern | Verifiable RL for Scientific Diagram Generation

July 2026 - Present

Lead author of **TexGemma-D1**, submitted to the August 2026 ACL Rolling Review cycle: an LLM training pipeline with a verifiable symbolic reward – the Diagram Semantic Verifier (DSV) – for scientific-diagram generation.

- ▷ Designed DSV, the verifiable symbolic reward at the core of TexGemma-D1: an 18-dimensional evaluator (23 with render checks) spanning 11 metric families that maps a TikZ program to a typed relational graph and scores what a diagram *asserts*, not how it looks – exact decision procedures where the semantics are decidable (Schwartz-Zippel genericity, DDAR-style geometric consistency, product-automaton search), explicit abstention where they are not, and a dual-modality witness that localizes each failing aspect to a source line/column *and* a pixel region of the render.
- ▷ Validated the metric without human labels: on 2,904 known-transform pairs from 728 real DaTikZ-V4 items, DSV separates semantics-preserving redraws from semantics-breaking edits at abstention-penalized AUROC 0.785 (95% CI [.769,.802]) versus 0.421 for a code-edit-distance baseline (*below* chance), records **0.0% strict inversions over 905,000+ pairwise comparisons** where DreamSim prefers a real corruption 35.1% of the time and code edit distance 56.1%, and beats two VLM judges (Gemini 2.5 Flash, Claude Sonnet 4.5) on the image side.
- ▷ Built the training pipeline: distilled explicit <think> reasoning traces into a Gemma-4-12B SFT checkpoint (LoRA, multi-GPU) so the policy plans structure before emitting LaTeX, then post-trained it with DAPO/Dr.GRPO against the DSV+RenderGraph verifiable reward behind a hard compile gate – raising DaTikZ-984 compilation from **61.3% to 91.4%**.
- ▷ Applied the same symbolic reward to a published Qwen2-based 3B DaTikZ SOTA policy (continued post-training): image similarity (DeTikZify-V2) rose from 0.260 to 0.353 and strict compilation from 93.8% to 98.8%, confirming the reward transfers across model families and lifts both correctness and validity.
- ▷ Enforced scientific rigor with matched controls and pre-registered ablations, reporting the ties and nulls: adding an R_{Sim} image-

similarity term to DSV+RenderGraph improved matched-description R_{Sim} (0.2674→0.2859) and compilation (94.7%→96.0%) but the gain *vanished* under prompt-distribution shift, and witness-gated per-token credit showed no measurable benefit – evidence that verifiable, distribution-robust reward, not perceptual similarity or handcrafted credit, is what drives the gains.

XtalPi Inc (QuantumPharm Inc) | XAB-Antibody IDD | Research Scientist Intern

Aug 2024 - July 2025

- ▷ Designed and implemented AI models for antibody binding affinity prediction. Reproduced training pipelines for binding-ddg-predictor and GearBind, developed data visualizations for the SKEMPI dataset, and established robust dataset splits to prevent leakage by ensuring complexes remain within the same fold. Conducted comparative analyses of focal loss (mutation-based), weighted MSE/BCE, balanced sampling, and differentiable ranking strategies.
- ▷ Developed a balanced sampling strategy for binding-ddg-predictor to ensure equal representation of mutations with and without alanine. Achieved performance improvements of approximately +10% in Pearson correlation and +5% in Spearman correlation.
- ▷ Integrated Carbondesign encoder into the binding-ddg-predictor framework to enhance encoder feature representations. Augmented the SKEMPIv2 dataset by generating multi-mutation variants, increasing sample size by ~10x. Successfully reproduced GearBind and conducted benchmarking experiments on the augmented SKEMPIv2 dataset, demonstrating improved model generalization.

Amazon Web Services | VMware Cloud on AWS (Brio) | Software Engineer Intern | Seattle, Wash- ington

May 2023 - Aug 2023

- ▷ Designed UI to automate the workflow to update Rate Card for Brio resource console using React, TypeScript and JavaScript.
- ▷ Integrated UI with backend API using Java, API Gateway and Amazon Lambda. Supported validation and creating and updating rate cards through csv. Preview visualizes modifications and identifies pricing errors. (5k+ lines of code in total).
- ▷ Designed anomaly detection algorithms for disbursement generator to identify usage spikes or subscription anomalies using Java, SQL, Amazon Athena, S3 Buckets, DynamoDB Table. Utilized CloudWatch to create tickets and alarms to engineers.

4Paradigm Co., Ltd | OpenMLDB | GitLink Code Camp 2022 | Open-Source Developer | Shanghai

July 2022 - Oct 2022

- ▷ Developed an automated feature engineering pipeline, including feature generation and selection, based on AutoX and OpenMLDB sql with Python (600+ lines of code). See pull request: <https://github.com/4paradigm/OpenMLDB/pull/2381>.
- ▷ Saving data scientists in dozens of companies' weeks of time by automatic feature selection pipeline.
- ▷ The data is transformed by OpenMLDB sql to get time series and statistics features, and then we select the most important K features based on algorithms like Adversarial Validation, GRN feature selection or Reinforcement Learning.

Intel Co., Ltd | DL Model Optimization Department | Deep Learning Software Engineer Intern | Shanghai

Nov 2021 - June 2022

- ▷ Implemented new features to [Intel® Neural Compressor](#), Cross Layer Equalization, a data free quantization to rescale different layers' weight range to reduce over 5% in the drop in accuracy after quantization like FP32 to INT8, and edited documents.
- ▷ Studied open-source software like Nvidia Triton Inference Server (around 70k lines of code) and AI Model Efficiency Tool and gave presentations to around 100 colleagues in whole department to introduce the design and technique detail of them.
- ▷ Cooperated to design and implement class of AI inference server software by C++, which enables the team to deploy trained AI models from multiple frameworks and deploy more models on GPU or CPU based infrastructure to simplify AI inferencing.

ShuKun Technology Co., Ltd | R & D Department | Algorithm Intern | Shanghai

Dec 2020 - Apr 2021

- ▷ Implemented multi-node and multi-GPU training with horovod framework, NVIDIA clara train sdk, OpenMPI, and NCCL2; 4GPU (2 nodes each 2GPU) can reach 2.5x speed up compared to 1GPU (Discovered that communication is bottleneck).
- ▷ Compared the efficiencies of models like 3D-UNet when multi-node training with different GPU configurations in Python and wrote the training process documents, saving algorithm scientists days in each training job (~0.65n times faster with n nodes).
- ▷ Used Java to add multi-machine and multi-GPU training functions to the company's back-end web page for algorithm scientists.

TEACHING EXPERIENCE

ICS 6B Boolean Logic & Discrete Structures | Professor: Wayne Hayes | Teaching Assistant

Winter 2025

University of California, Irvine.

ICS 6D Discrete Mathematics | Professor: Jing Zhang | Teaching Assistant

Spring 2025

University of California, Irvine.

VE370 Intro to Computer Organizations | Professor: Gang Zheng | Teaching Assistant

Sep 2021 - Dec 2021

University of Michigan - Shanghai Jiao Tong University Joint Institute.

PERSONAL PROJECTS

AgenticVBench: Long-Horizon Video-Agent Tasks

Jul 2026 - Present

- ▷ Shipped a 238-minute Minecraft action-ledger benchmark (1,995 timed events) with machine-exact ground truth, deterministic order/time verification, anti-shortcut tests, and three-agent calibration [[Issue #74](#); [merged PR #108](#)].
- ▷ Built a 55-minute, 12-race SuperTuxKart telemetry task with reproducible generation, machine-exact ground truth, coverage-hardened scoring, and multi-agent calibration; under upstream review [[Issue #73](#); [PR #106](#)].

Photon 47 – Browser WebGL Game with a Serverless Postgres Backend (Solo)

2026

TypeScript/JavaScript, WebGL, PostgreSQL (Supabase), Deno Edge Functions, Python (Playwright); shipped to 5 platforms.

- ▷ Designed a 13-table PostgreSQL/Supabase backend with 38 SECURITY DEFINER stored procedures (profiles, leaderboards, ELO matchmaking, payments, analytics), secured with Row-Level Security and idempotent migrations; concurrency-safe matchmaking via `SELECT ... FOR UPDATE SKIP LOCKED` and a transactional ELO update (K=32).

- ▷ Built an event-analytics layer – append-only JSONB event ingestion plus time-windowed SQL rollups (GROUP BY / COUNT(*) FILTER) for mode popularity, character pick rates, item sales, and per-mode win rates.
- ▷ Shipped a Stripe payments pipeline as Deno/TypeScript Edge Functions with JWT/JWKS-verified checkout and an idempotent, signature-verified fulfillment webhook, deployed via an env-driven test/live script.
- ▷ Built a 95-suite Playwright harness (~6,850 lines) and one-source release pipeline for web, Steam, CrazyGames, Poki, and itch.

COURSE PROJECTS

CS 425 Distributed Systems: Distributed ML Inference System (C++ and Python) (100/100) Aug 2022 - Dec 2022

- ▷ Implemented scheduling algorithm to obtain fair-time inference, making each ML job query rate within 20% difference.
- ▷ Design a distributed file system to maintain datasets with SWIM-like membership protocol (nearly 1e-2 false positive rate).
- ▷ Supported fault-tolerance and can recover when N members and coordinator leave or fail. Leader re-election is within 10s.

CS 543 Computer Vision (A): (Python) Aug 2022 - Dec 2022

- ▷ Channel Alignment: Designed multi-scale algorithm to align 3 channels of images, evaluating with NCC and Fourier Transform.
- ▷ Laplacian Blob Detection: Implemented a scale-space blob detection with a Laplacian scale space using Laplacian of Gaussian filter and performed non-maximum suppression in scale space.
- ▷ Stitching: Stitching pairs of images using OpenCV, SIFT descriptors and RANSAC to estimate homography mapping of images.
- ▷ 3D Shape from shading: Estimate the albedo and surface normal and compute the surface height map by integration to reconstruct 3D surface from image of shading.

CS 438 Communication Network (A+): TCP Protocol (C++) Aug 2022 - Dec 2022

- ▷ Implemented transport protocol with properties equivalent to TCP based on unreliable UDP, which can tolerate packet drops, allow other concurrent connections fair chance (0.5x to 2x of TCP), doesn't give up entire bandwidth to other connections.

Sonar detection water pipe state system based on machine vision and neural network compression Sep 2021 - Dec 2021

ECE4710 Advanced Embedded Systems (MDE, Thesis). | Instructor: Zou An, Assistant Professor of UM-SJTU Joint Institute.

- ▷ Processed the source data of sonar as a matrix and use Python to generate polar images in real time.
- ▷ Used OpenCV to perform subtraction and corrosion operations to complete hole or bulge detection.
- ▷ Designed a compressed and binarized YOLO architecture to realize defect area detection with accuracy 99%.
- ▷ Used HLS to deploy to FPGA, matrix parallel operation, real-time embedded system is implemented.

HONORS & AWARDS

[Kaggle ARC Prize 2026: ARC-AGI-3](#) – 273rd of 3,105 ranked teams (top 8.8%; team “Ryan #3,” account “explcre”; public score 3.53; within the current public leaderboard’s bronze-medal range, Sep 17, 2026) 2026

Shanghai Jiao Tong University Outstanding Graduate (School Level) Aug 2022

Microsoft Imagine Cup 2021, China Regional Final – Third Prize Jan 2021

Mathematical Contest in Modeling – Meritorious Winner, team 2006402 [[official results](#)] Apr 2020

“Jidong Cup” CCVR China Virtual Reality Competition, Future Virtual Reality Product Group – Second Prize Sep 2020

Academic Year Undergraduate Excellence Scholarship (2018–2019 & 2020–2021) Nov 2019 / 2021

SKILLS

- ▷ Programming: C/C++ (Proficient), Python (Proficient), MATLAB, R, Verilog, Java, RISC-V Assembly, SQL, Shell, TypeScript, CUDA Programming.
- ▷ Frameworks & Libraries: PyTorch, TensorFlow, horovod, Keras, Sk-learn, Pandas, NumPy, OpenCV, Matplotlib, Selenium.
- ▷ Developer Tools: Docker, Git, LaTeX, PS/Spice, Linux.
- ▷ Languages: English (Fluent, Full Professional Proficiency), Chinese (Native Speaker), German (Beginner).