

EDUCATION

- Stanford University Sep 2018 - Jun 2022
- B.S., Computer Science (AI Specialization) • GPA: 4.18 / 4.0 • Graduated with Distinction (Summa Cum Laude)
 - [Endorsed](#) by Fields Medalist Prof. [Ngo Bao Chau](#) for exceptional mathematical talent and research potential.

AWARDS

- | | |
|---|---|
| <ul style="list-style-type: none"> International Mathematical Olympiad: Bronze Medal Vietnam Math Olympiad: Gold Medal (Rank #1) Hanoi Open Math Competition: Gold Medal | <ul style="list-style-type: none"> American Mathematics Competition: Gold Medal International Math Local Tournament: Gold Medal Regional Math Competitions: 5+ Gold Medals |
|---|---|

WORK EXPERIENCE

- OpenAI Member of Technical Staff, Researcher Feb 2026 - Present
- Research novel end-to-end **LLM post-training, reinforcement learning, and test-time compute scaling** strategies to unlock breakthrough capabilities in state-of-the-art foundation models and production-grade agents.
 - Architect cluster-scale algorithms, dynamic RL environments, high-throughput pipelines, robust recipes, and evals to redefine the frontier of **steerability, instruction following, reasoning, planning, and agentic** capabilities.

- Meta Senior Machine Learning Engineer Jul 2022 - Feb 2026
- **Ranked #1** for organizational impacts (Exceeded Expectations+), delivering **90-120%** of revenue targets and **40-50%** of DAU goals through development and ownership of high-leverage ML solutions.
 - Architected end-to-end **post-training and serving pipelines for compact LLMs**, optimizing **fine-tuning and alignment** for high-scale generation and ranking. Engineered robust **LLM-based guardrails** to preemptively neutralize safety, integrity, and quality risks, enforcing a critical pre-deployment defense layer.
 - Pioneered next-generation modeling architectures by developing novel **agentic recommendation, reinforcement learning, and causal inference** technologies for multi-objective and long-horizon user value optimization. Advanced model capabilities through **knowledge distillation, Transformer-based sequence modeling, and contextual and multimodal representation learning**, capturing deep latent interests to drive engagement.
 - Led **large-scale, cross-organizational initiatives spanning 20+ engineers** across infrastructure, product, and modeling teams. Independently defined strategies, aligned stakeholders, and executed system-wide upgrades that materially advanced Meta's notification recommendation system.
 - **Scaled team capability by 3x** by proactively scoping and launching new technical roadmaps, mentoring and onboarding 10+ IC4/IC5 engineers, and enabling sustained growth and research-grade innovation.

- Meta Software Engineering Intern Jun 2021 - Sep 2021
- Developed large-scale ML models in PyTorch to generate **contextual embeddings** of users and ads, enhancing semantic understanding and fine-grained personalization in ad retrieval and ranking systems.
 - Designed and deployed **representation learning frameworks** for user-ad matching across multiple surfaces, significantly improving relevance and driving **ad revenue growth**. Used Python, SQL, C++, PHP, and Spark.

- Meta Software Engineering Intern Jun 2020 - Sep 2020
- Enhanced latency profiling tools with module-level debugging. Achieved a **5x increase in runtime efficiency**.
 - Accelerated Conv1D and channel shuffle operations via **low-level kernel optimizations for PyTorch**, delivering up to **10x operator-level speedup** for on-device speech and NLP models. Used Python, C, and C++.

RESEARCH EXPERIENCE

- Stanford AI Lab Undergraduate Researcher Apr 2020 - Sep 2021
- Built large-scale data pipelines in Python, NumPy, and Pandas to ingest, align, and preprocess satellite imagery and global forest loss driver labels. Implemented advanced **data augmentation** and **stratified sampling** techniques to improve signal diversity and model robustness across diverse geographies.
 - Designed and trained deep learning models in PyTorch, including **CNNs, LSTMs, and multimodal fusion architectures**, to classify forest loss drivers from multi-temporal satellite imagery, achieving **80% classification accuracy** and supporting scalable environment monitoring.

- Computer Science Lab Undergraduate Researcher Sep 2019 - Dec 2019
- Researched model compression techniques, including **structured pruning, regularization-based sparsity, and weight quantization**, to reduce the computational footprint of deep convolutional networks.
 - Applied regularization and pruning strategies to ResNet, achieving a **15% reduction in FLOPs** with a **2% improvement in accuracy**, demonstrating efficient compression without performance trade-offs.

PROJECTS

- Nano-Reasoner: Unified Post-Training Framework for Math Reasoning Models [\[code\]](#)
- Designed an end-to-end research-grade, memory-efficient **post-training** pipeline for **reasoning models** with **SFT** and **reinforcement learning** in PyTorch. Implemented cutting-edge reinforcement learning algorithms (**PPO**, **GRPO**, **Dr. GRPO**, **GSPO**, **DAPO**), achieving **20% accuracy wins** in math reasoning benchmarks.
- PRM-Math: Inference-Time Compute Scaling via Dense Process Supervision [\[code\]](#)
- Fine-tuned a **Process Reward Model** in PyTorch to assign dense, step-wise scores to reasoning trajectories.
 - Implemented cutting-edge **test-time compute scaling** strategies (**Best-of-N Search**, **PRM-Weighted Voting**, and **Monte Carlo Tree Search**), resulting in a **10% accuracy improvement** in math reasoning benchmarks.
- Nano-Video-Gen: Spacetime Diffusion Transformer with Rectified Flow Matching [\[code\]](#)
- Engineered a **Diffusion Transformer video generator** from scratch in PyTorch with **3D spacetime patch embeddings** and **Rectified Flow Matching**. Achieved convergence on Moving MNIST within 48 hours on a single GPU by optimizing the Flash Attention mechanism and implementing mixed-precision training.
- Agentic Deep Research: Recursive Reasoning and Self-Correction Engine [\[code\]](#)
- Designed a **recursive research agent** in LangGraph that executes **dynamic DAGs** for complex, multi-hop query resolution and utilizes **cyclic self-correction loops** to refine search queries. Built a **multi-agent orchestration layer** where sub-agents parallelize retrieval, synthesis, and critique, significantly reducing hallucination rates.
 - Developed a rigorous **LLM-as-a-Judge evaluation** pipeline, demonstrating improvements in **perfect scores (+14%)**, **factuality (+14%)**, **completeness (+3%)**, and **coherence (+2%)** over naive RAG baselines.
- Video-DPO: Temporal Alignment for Video Diffusion via Direct Preference Optimization [\[code\]](#)
- Implemented a **synthetic data** pipeline and adapted **DPO for video diffusion** models to align generation for temporal stability and motion smoothness, **reducing temporal warping artifacts by 25%**.
- Visual-CoT: Pixel-Grounded Reasoning with Multi-Modal Chain-of-Thought [\[code\]](#)
- Fine-tuned a **Vision-Language Model** to generate **interleaved reasoning traces** containing both text and precise bounding box coordinates, forcing the model to spatially attend to evidence. Demonstrated a significant **reduction in object hallucination** and improved **interpretability on complex visual reasoning tasks**.
- Tool-Use DPO: Schema-Constrained Alignment via Identity Preference Optimization [\[code\]](#)
- Fine-tuned an **LLM for rigid API contract adherence** using **IPO**. Engineered a **synthetic dataset** to penalize subtle schema violations, **reducing syntax errors by 16%** on complex nested JSON generation tasks.
- Web-Browser-Agent: Multimodal Autonomous Web Navigation via Visual Grounding [\[code\]](#)
- Engineered a **multi-modal autonomous web agent** utilizing a **Vision-Language Model** and a hybrid **Set-of-Mark** pipeline to execute long-horizon browser workflows. Architected a **Verify-Act-Verify** reasoning loop with strict pixel-difference evaluation to detect latent execution failures and trigger self-correction.
- Efficient-Reasoner: Adaptive Compute Allocation via Reinforcement Learning [\[code\]](#)
- Fine-tuned an **LLM** via **SFT** and **GRPO** to autonomously route between **System 1 (fast zero-shot)** and **System 2 (tool-augmented)** reasoning. Engineered a **Value of Information reward function** that balances multi-hop QA correctness, strict XML schema adherence, and computational cost. Maintained peak SFT empirical accuracy while **reducing inference token consumption by 40%** and **redundant API calls by 60%**.

PUBLICATIONS

- FlapAI Bird: Training an Agent to Play Flappy Bird Using Reinforcement Learning [\[paper\]](#) [\[code\]](#)
- How Not to Give a FLOP: Combining Regularization and Pruning for Efficient Inference [\[paper\]](#) [\[code\]](#)
- Privacy Preserving Inference of Personalized Content for Out of Matrix Users [\[paper\]](#) [\[code\]](#)
- GANime: Generating Anime and Manga Character Drawings from Sketches [\[paper\]](#) [\[code\]](#)
- BERT-VQA: Visual Question Answering on Plots [\[paper\]](#) [\[code\]](#)
- Pixel-Perfect Piloting: Superhuman Control of Pixelcopter via Reinforcement Learning [\[paper\]](#) [\[code\]](#)

SKILLS

Languages & Technologies: Python • C • C++ • SQL • JavaScript • PyTorch • JAX • TensorFlow • vLLM
HuggingFace • TRL • NumPy • Pandas • Accelerate • FSDP • LangGraph • PrimeRL • SkyRL • Scikit-learn
Unsloth • Spark • React Native • HTML • CSS • Express.js • Node.js • React • MongoDB • Next.js • Django
Coursework: Machine Learning • Deep Learning • LLMs • NLP • Reinforcement Learning • Computer Vision
Information Retrieval • Spoken Language Processing • Graph ML • Parallel Computing • Data Structures &
Algorithms • Systems • Web Applications • Statistics • Linear Algebra • Convex Optimization • Game Theory