

# Minghui Liu

minghui@umd.edu | [GitHub](#) | [Webpage](#) | [Google Scholar](#)

Looking for: AI/Machine Learning Research Internship | Starting Summer 2026.

Research Interests: LLM Reasoning, RL & Curriculum Learning, Self-Improvement, KV Cache Compression

## EDUCATION

---

### University of Maryland, College Park

Jan. 2021 – Present

*Ph.D. Candidate, Computer Science.*

*Advisor: Prof. Furong Huang*

### Trinity College, Hartford CT

Sep. 2013 – May. 2017

*B.Sc in Computer Science. summa cum laude with distinction.*

*Top 3%*

*Phi Beta Kappa. Dean's Scholar. Faculty Honor.*

*Top 5%*

## ON GOING PROJECTS

---

### 1. Ladders-RL: Dynamic Difficulty Adjustment for Reinforcement-Trained Reasoning

Extends GRPO by dynamically switching to easier progressive rewrites of the same question when all rollouts in a group fail to receive reward. This on-the-fly fallback preserves gradient signal under sparse-reward settings, enabling stable reinforcement learning on reasoning tasks where baseline GRPO collapses due to zero-reward episodes.

### 2. PRM-GRPO: Process-Reward-Model Guided Reinforcement for Reasoning Chains

Investigates the complementary strengths of Process Reward Models (PRMs) and outcome-verifier rewards in reinforcement training of reasoning chains. Demonstrates that PRMs yield smoother optimization and better step-level alignment under sparse-reward regimes, while outcome verifiers enforce global logical validity. Proposes a unified training framework and derives practical guidelines on when and how to deploy PRMs effectively for reasoning tasks

## SELECTED RESEARCH & PUBLICATIONS (2025–2026)

---

### 1. Ladders of Thought: A Self-Evolving Curriculum of Progressively Simplified Reasoning Traces

ICLR 2026 Submission

*M. Liu, T. Magelinski, D. Yuan, Q. Yu, F. Huang*

- Developed a progressive rewrite framework that generates semantically faithful but easier variants of reasoning tasks.
- Introduced a self-evolving bandit scheduler that adapts task difficulty to model learning progress.
- Achieved up to +32pp improvement in arithmetic reasoning and +25pp in multi-hop reasoning tasks for small LLMs (OPT/Pythia).
- Demonstrated faster convergence and better generalization across datasets, highlighting the benefits of adaptive curricula.

### 2. Hold Onto That Thought: Assessing KV Cache Compression on Reasoning

ICLR 2026 Submission

*M. Liu, T. Rabbani, A. Palnitkar, K. R. Sang, H. Jae, D. Yao, S. Shabihi, F. Zhao, T. Li, C. Zhang, K. Zhang, F. Huang*

- Benchmarked over 25 KV cache compression strategies for long-context reasoning tasks (GSM8K, MATH-500, FOLIO, StrategyQA).
- Integrated decoding-phase compression into NVIDIA's `kvpress` and analyzed reasoning degradation patterns under varying cache budgets.
- Provided new evaluation metrics quantifying reasoning stability versus compression ratio and sequence length.

### 3. MORSE: A Suite of Programmatically Controllable Multimodal Reasoning Environments with Steerable Difficulty

[arXiv 2025](#), ICLR 2026 Submission

*Z. Cai, A. Wang, A. Satheesh, A. Nakhawa, H. Jae, K. Powell, M. Liu, N. Jay, S. Oh, X. Wang, Y. Liang, T. Goldstein, F. Huang*

- Built MORSE, a programmatic video reasoning simulator enabling controllable spatial, temporal, and abstract task difficulty.

- Released MORSE-500, a multimodal benchmark of 500 reasoning tasks covering planning, spatial awareness, and object relations.
  - Implemented an automated generator using LLM-in-the-loop feedback to refine task diversity and correctness.
4. **PersonaLedger: Generating Realistic Financial Transactions with Persona-Conditioned LLMs and Rule-Grounded Feedback** ICLR 2026 Submission  
*D. Yuan, X. Liu, M. Liu, T. Farnan, S. Tesliuc, D. L. Bergman, J. Montgomery, N. H. Nguyen, C. B. Bruss, F. Huang*
- Proposed a hybrid LLM–programmatic engine for generating persona-grounded synthetic financial transactions.
  - Built a dataset of 30M transactions across 23K users with strict accounting constraints, rule enforcement, and regeneration logs.
  - Released two benchmark tasks—illiquidity prediction and anomaly detection—supporting reproducible financial AI research.
5. **HashEvict: Towards Pareto-Efficient KV Cache Eviction via Locality-Sensitive Hashing** [arXiv 2025](#)  
*M. Liu, T. Rabbani, T. O’Halloran, A. Sankaralingam, C. Fermüller, Y. Aloimonos, M. A. Hartley, F. Huang*
- Designed an adaptive KV cache compression algorithm leveraging SimHash-based cosine similarity to identify token redundancy pre-attention.
  - Achieved 30–70% cache reduction and up to 2× prefill and decoding speedups across reasoning, summarization, and retrieval tasks.
  - Proposed a Pareto-efficiency metric to balance speed and language quality, formalizing trade-offs in KV cache management.
  - Conducted large-scale benchmarking on LLaMA and Mistral models; released open-source code and ablation results.

WORK EXPERIENCE

---

**Applied Research Intern**  
*Capital One, New York NY*

June 2025 – Aug 2025

- Contributed to the *Ladders-of-Thought* project; designed adaptive curricula and evaluated reasoning improvements leading to ICLR 2026 submission.
- Developed adaptive reasoning curricula for small LLMs, integrating progressive simplification and dynamic difficulty control.
- Designed experiments demonstrating significant reasoning gains and convergence speed improvements.

**HPC Software Engineer**  
*Altair Engineering, Sunnyvale CA*

Aug 2017 – Jan 2021

[PROJECT](#) [GITHUB](#)

- Enhanced OpenPBS job scheduler accuracy, added Docker support, and optimized Linux build process.

**Research Assistant**  
*Trinity College, Hartford CT*

Jan 2014 – Sep 2016

*Topic: Hypergraph-based Image Matching Algorithm on the GPU*

- Developed GPU-based hypergraph image matching achieving 8–20× speedup; [presented](#) at NVIDIA GPU Technology Conference.

RELEVANT COURSES

---

<i>Fundamentals</i>	OOP, Discrete Math, Data Structures and Algorithms
<i>Artificial Intelligence</i>	Machine Learning, Computer Vision, Natural Language Processing
<i>Algorithms</i>	Advanced Algorithm Design, Computational Geometry
<i>Systems</i>	Operating Systems, Database, High Performance Computing
<i>Mathematics</i>	Linear Algebra, Calculus, Statistics, Probability, Numerical Analysis, Graph Theory, Combinatorics

SKILLS

---

<i>Python - Machine Learning</i>	pytorch, huggingface, numpy, pandas, scipy, matplotlib, sklearn, web-crawling (bs4), ...
<i>High Performance Computing</i>	C, CUDA, cython, MPI, OpenMP
<i>Front-End, Back-End</i>	HTML, CSS, Javascript