

SHOUREN WANG

☎ +1 (929) 302-1790 ✉ sxw992@case.edu 🌐 [Personal Website](#)
🌐 [LinkedIn](#) 🐙 [GitHub](#) 🎓 [Google Scholar](#)

EDUCATION

Case Western Reserve University <i>Ph.D. in Computer Science</i>	Cleveland, OH <i>Aug. 2024 – Present</i>
New York University <i>M.S. in Computer Engineering GPA: 3.93/4.0</i>	Brooklyn, NY <i>Sep. 2021 – May. 2023</i>
Hunan University <i>Undergraduate in Software Engineering</i> <i>Undergraduate and B.E. in Digital Media Technology GPA: 83.16/100.0</i> <i>(A specialization of Software Engineering focusing on visual software/algorithm development)</i>	Changsha, China <i>Sep. 2015 – June 2017</i> <i>June. 2017 – June 2019</i>

EXPERIENCE

Research Intern <i>NYU Game Innovation Lab</i>	Aug. 2023 – July 2024 <i>Brooklyn, NY</i>
• Explored the research and methods for creative AI for video games. • Developed enjoyment focused game-playing agents for Street Fighter 2 via an extended PPO model and self-play. • Advisor: Julian Togelius, Professor, New York University.	
Research and Development Engineer <i>AsiaInfo Technology</i>	Sep 2019 – Sep 2020 <i>Nanjing, China</i>
• Tested and enhanced the performance of CTDI project as a QA group member, contributed to project efficiency. • Manager: Ted Wang.	
Teaching Assistant <i>Case Western Reserve University</i>	Sep. 2024 – May 2026 <i>Cleveland, OH</i>
• Teaching assistant for <i>Large Language Models</i> and <i>Database Systems</i>: graded assignments and held office hours. • Professors: Xiaotian Han (Large Language Models); Yinghui Wu (Database Systems).	

PROJECTS

Agent Team Work Zone [GitHub] <i>LLM Agents, Claude Code, Multi-Agent Workflow</i>	Apr. 2026 – Present
• Built and maintain an open-source persistence layer for Claude Code’s Agent Teams, whose teammates vanish on session restart with no native recovery. • Designed per-agent workstations (role, working context, checkpoints, commitments) and a team registry, so one command rebuilds an interrupted team from disk. • Shipped 13 slash-command skills and hook-driven auto-checkpointing across four public releases, maintaining parallel editions through a breaking Claude Code API change. • Dogfood the layer end-to-end: the agent team that develops it runs on it.	
Path-Lock Expert: Efficient Hybrid Thinking Architecture [Paper] <i>LLMs, Reasoning</i>	Jan. 2026 – Jul. 2026
• <i>Accepted at COLM 2026.</i> • Developed a parameter-efficient architecture approach to improve hybrid-thinking <i>mode controllability</i> (thinking vs. fast) without compromising overall task performance. Built an efficient training pipeline with curated data. • Experiments on reasoning benchmarks (AIME24, MATH500, GPQA-Diamond) indicate improved mode separation/controllability with competitive accuracy, and greatly improved accuracy for no-think mode.	
Demystifying Hybrid Thinking [Paper] <i>LLMs, Reasoning, Post-Training</i>	Sep. 2025 – Feb 2026
• <i>Accepted at EMNLP 2026 Main.</i> • Systematically analyzed the controllability of hybrid-thinking LLMs (think vs. no-think), showing incomplete mode separation and quantifying “reasoning leakage” via output length and reflection-token counts across multiple evaluations (MATH500/AIME24/GPQA/MMLU-STEM). • Conducted controlled SFT ablations on data scale (20k–140k), pairing strategy, think:no-think ratio, and two-phase training; proposed a practical training recipe that preserves accuracy while cutting no-think verbosity (1085→585 tokens on MATH500) and reflection-token frequency (5917→522).	

- *Accepted at ACL 2026 Findings.*
- Identified via attention analysis that hybrid-thinking mode switching is dominated by a few *token-level triggers*—“Okay” activates reasoning while `</think>\n\n` suppresses it—rather than high-level Think/No-think instructions.
- Proposed **Mid-Think**, a training-free prompting format combining both triggers to achieve Pareto-optimal accuracy–length trade-offs (MATH500: 92.1% acc. at 2589 tokens vs. 94.4% at 4904), outperforming fixed-token and prompt-based baselines across Qwen3-4B/8B/14B/32B.
- Applied Mid-Think to GRPO-based RL after SFT, cutting training time $\sim 15\%$ (54h $\rightarrow$ 46h) while improving Qwen3-8B on AIME (69.8% $\rightarrow$ 72.4%) and GPQA (58.5% $\rightarrow$ 61.1%).

PUBLICATIONS

*Equal contribution. †Corresponding author(s).

[1] **Path-Lock Expert: Separating Reasoning Mode in Hybrid Thinking via Architecture-Level Separation**

COLM 2026

- Authors: **Shouren Wang***, Wang Yang*, Chuang Ma, Debargha Ganguly, Vikash Singh, Chaoda Song, Xinpeng Li, Xianxuan Long, Vipin Chaudhary†, Xiaotian Han†
- Links: [\[Paper\]](#), [\[Repository\]](#)

[2] **Demystifying Hybrid Thinking: Can LLMs Truly Switch Between Think and No-Think?**

EMNLP 2026 Main

- Authors: **Shouren Wang***, Wang Yang*, Xianxuan Long, Qifan Wang, Vipin Chaudhary†, Xiaotian Han†
- Links: [\[Paper\]](#), [\[Repository\]](#)

[3] **Mid-Think: Training-Free Intermediate-Budget Reasoning via Token-Level Triggers**

ACL 2026 Findings

- Authors: Wang Yang*, **Shouren Wang***, Debargha Ganguly, Xinpeng Li, Chaoda Song, Vikash Singh, Vipin Chaudhary†, Xiaotian Han†
- Links: [\[Paper\]](#), [\[Repository\]](#)

[4] **A Survey on Efficient Protein Language Models**

TMLR 2026

- Authors: **Shouren Wang**, Debargha Ganguly, Vinooth Rao Kulkarni, Wang Yang, Zhuoran Qiao, Daniel Blankenberg, Vipin Chaudhary†, Xiaotian Han†
- Links: [\[Paper\]](#), [\[Repository\]](#)

[5] **When Domains Interact: Asymmetric and Order-Sensitive Cross-Domain Effects in Reinforcement Learning for Reasoning**

Under Review

- Authors: Wang Yang, **Shouren Wang**, Chaoda Song, Chuang Ma, Xinpeng Li, Nengbo Wang, Kaixiong Zhou, Vipin Chaudhary†, Xiaotian Han†
- Links: [\[Paper\]](#), [\[Repository\]](#)

[6] **A Survey on Agent Skills for LLMs: A Lifecycle Perspective from Construction to Ecosystems**

Under Review

- Authors: Wang Yang, Chaoda Song, Xinpeng Li, **Shouren Wang**, Nengbo Wang, Yanyan Zhang, Chuang Ma, Debargha Ganguly, Vikash Singh, Shuai Xu, Jing Ma, Yu Yin, Vipin Chaudhary†, Xiaotian Han†
- Links: [\[Paper\]](#), [\[Repository\]](#)

[7] **AgentCE-Bench: Agent Configurable Evaluation with Scalable Horizons and Controllable Difficulty under Lightweight Environments**

Under Review

- Authors: Wang Yang, Chaoda Song, Xinpeng Li, Debargha Ganguly, Chuang Ma, **Shouren Wang**, Zhihao Dou, Yuli Zhou, Vipin Chaudhary†, Xiaotian Han†
- Links: [\[Paper\]](#), [\[Repository\]](#)

[8] **MemTrace: Probing What Final Accuracy Misses in Long-Term Memory**

Under Review

- Authors: Xianxuan Long, Zhikai Chen, Shenglai Zeng, **Shouren Wang**, Kai Guo, Jiliang Tang[†]
- Links: [\[Paper\]](#)

[9] **Reliability-Gated Source Anchoring for Continual Test-Time Adaptation**

Under Review

- Authors: Vikash Singh, Debargha Ganguly, Weicon Chen, Sabyasachi Sahoo, Sreehari Sankar, Biyou Zhang, Mohsen Hariri, **Shouren Wang**, Osama Zafar, Christian Gagné, Vipin Chaudhary[†]
- Links: [\[Paper\]](#)

[10] **CausalGuard: Conformal Inference under Graph Uncertainty**

Under Review

- Authors: Vikash Singh, Weicon Chen, Debargha Ganguly, Yanyan Zhang, Nengbo Wang, Sreehari Sankar, Mohsen Hariri, Alexander Nemecek, Chaoda Song, **Shouren Wang**, Biyou Zhang, Van Yang, Erman Ayday, Jing Ma, Vipin Chaudhary[†]
- Links: [\[Paper\]](#)

[11] **Enhancing Player Enjoyment with a Two-Tier DRL and LLM-Based Agent System for Fighting Games**

Under Review

- Authors: **Shouren Wang**, Zehua Jiang, Fernando Sliva, Sam Earle, Julian Togelius[†]
- Links: [\[Paper\]](#), [\[Repository\]](#)

ACADEMIC SERVICE

Reviewer: Conference on Language Modeling (COLM) 2026

TECHNICAL SKILLS

Languages: Python, Java, MATLAB

Frameworks: PyTorch, Transformers, vLLM, Llama-Factory, Stable-Baselines3, OpenAI Gym, Scikit-Learn

Developer Tools: Git, Slurm, Anaconda, Docker, Singularity, Linux, VS Code

AI Development Tools: Claude Code, OpenClaw, Codex, Cursor

RELEVANT GRADUATE-LEVEL COURSEWORK

- **LLMs & Generative Models:** Large Language Models [CWRU, CSDS 600], Deep Generative Models [CWRU, CSDS 570]
- **Machine Learning:** Advanced Machine Learning [NYU, ECE-GY 7143], Machine Learning [NYU, ECE-GY 6143], Causality and Machine Learning [CWRU, CSDS 452], Data Mining [CWRU, CSDS 435]
- **Core CS & Mathematics:** Data Structures and Algorithms [NYU, ECE-GY 9343], Computing Systems Architecture [NYU, ECE-GY 6913], Probability and Stochastic Processes [NYU, ECE-GY 6303]