Dehao Wu

Email: dehaowu2@illinois.edu | jacob20030316-crypto.github.io

RESEARCH INTEREST

Agentic AI and Large Language Models, with a focus on Multi-Agent Systems, AI for Scientific Discovery, Biomedical AI, and Multimodal Learning.

EDUCATION

University of Illinois Urbana-Champaign (UIUC) Urbana, IL
M.S. in Information Management, GPA: 4.00/4.00 Expected in May 2027
Honors & Awards: Gorecki Endowed Scholarship (2,000 USD)

Southwest Petroleum University Sichuan, China
B.E. in Artificial Intelligence, GPA: 4.01/5.00, Ranking: 1/77 Sep. 2021 – June 2025
Honors & Awards: Provincial Outstanding Graduate Award, 2025; First-Class Academic Scholarship, 2021-2022, 2022-2023, 2023-2024 & 2024-2025

University of Texas at Austin Austin, TX
Summer Academic Program, GPA: 4.00/4.00 July 2024 – Aug. 2024

SELECTED RESEARCH PROJECTS

Agentic AI Uncovers Conserved Cross-Tissue Protein Co-Abundance Programs Inaccessible to Single-Dataset Analysis
Co-First Author | Supervisor: Haohan Wang | Under Review at *Bioinformatics* Jan. 2026 – Present

- Co-designed and co-developed a multi-agent framework for cross-tissue proteomics analysis, contributing to data analysis, biological interpretation, and manuscript preparation.
- Integrated protein expression data from ProteomicsDB to construct protein co-abundance networks across 41 human tissues and biofluids, systematically analyzing 820 tissue pairs.
- Constructed consensus networks and applied Leiden community detection, permutation testing, and multi-source biomedical evidence integration, identifying 1,833 conserved protein co-abundance clusters across 406 tissue pairs.
- Revealed that the breadth of tissue-level connectivity and the depth of pairwise co-abundance conservation represent distinct dimensions of cross-tissue organization that cannot be captured by a single connectivity ranking.
- Established a systematic resource for investigating cross-organ disease mechanisms and potential therapeutic targets through large-scale, agent-driven proteomics analysis.

Large-Scale Agentic Analysis Identifies Recurrent Regulators Linking Diseases across Patient Subgroups
First Author | Supervisor: Haohan Wang | Under Review at *npj Systems Biology and Applications* Mar. 2026 – Present

- Developed an agentic framework for large-scale analysis of heterogeneous GEO and TCGA transcriptomic cohorts, including metadata parsing, clinical-variable normalization, and patient stratification.
- Standardized patient characteristics into 86 clinical axes and constructed subgroup-specific gene regulatory networks, forming a network of gene regulatory networks across diseases and patient populations.
- Analyzed 132 diseases, 911 cohorts, and 150,000+ samples, performing more than 210,000 patient-subgroup network comparisons and identifying 222 cross-disease relationships with stable regulators and reproducible regulatory programs.
- Developed a reproducibility-oriented framework that distinguishes network-level similarity from regulator-level recurrence and ranking reproducibility, enabling more reliable identification of cross-disease regulatory mechanisms and candidates for functional validation and drug repurposing.

Large-Scale Repository Engineering via Agent-Native Reusable Code Primitives
Fifth Author | Supervisor: Haohan Wang | Under Review at *ICLR 2027* May–Sep. 2026

- Contributed to LEGO, an agent framework that constructs software repositories by retrieving and adapting reusable code primitives, integrating task-specific code, and revising implementations through testing.

- Contributed to the development of CodeFace, a searchable library of 1,424 validated code primitives with implementations, interface contracts, dependencies, tests, and provenance.
- The framework improved delivery scores by an average of 0.1474 across 13 models on 522 repository reconstruction tasks and achieved 14.9%–17.6% performance gains over matched baselines across three external benchmarks.

Composing Task-Specific Agent Harnesses at Test Time with Reusable Primitives

Third Author | Supervisor: Haohan Wang | Under Review at *ICLR 2027* May–Sep. 2026

- Contributed to STITCH, a framework that uses LLM-based primitive selection and deterministic, contract-guided compilation to construct task-specific agent harnesses without generating new mechanism code at test time.
- Contributed to developing reusable Harness Primitives from recurring agent failures, incorporating mechanism correctness and task-level benefit checks, explicit application scopes, and composition contracts.
- The framework achieved 80.5% and 72.2% Pass@1 on sampled SWE-bench Verified (100 tasks) and Terminal-Bench 2 (45 tasks), outperforming MetaHarness by 7.0 and 16.6 percentage points and Codex CLI by 1.5 and 15.5 points, respectively.

Focus on What You Need: Global-to-Local Dynamic Attention for Multimodal Sarcasm Detection

First Author | Supervisor: Chengxiang Zhai | Under Review at *ICASSP 2027* Mar. 2026 – Present

- Proposed Global-to-Local Dynamic Attention (GLDA) to address interference from irrelevant visual objects and background information in multimodal sarcasm understanding.
- Developed a global multimodal fusion module to capture image-text semantic incongruity, followed by local gated cross-attention to dynamically select visually and linguistically relevant regions and tokens.
- Designed a four-branch optimization framework integrating text, image, global fusion, and local gated representations to jointly improve unimodal and cross-modal representations.
- Achieved competitive performance on MMSD and MMSD2.0 while substantially reducing training time and GPU memory consumption compared with supervised fine-tuning of large vision-language models.
- Demonstrated cross-task generalizability by applying GLDA to multimodal rumor detection on English and Chinese datasets.

Agents' Last Exam (ALE)

Data Contributor | NeurIPS 2026 Apr. – May 2026

- Contributed to the development of a benchmark for evaluating AI agents on long-horizon professional workflows, with a focus on planning, tool use, and multi-stage task execution.
- Designed and implemented two long-horizon evaluation tasks with deterministic verification mechanisms as a Data Contributor.
- Developed a TCGA breast cancer differential-expression analysis task, including data packages, literature-supported ground truth, task specifications covering preprocessing, statistical testing, multiple-testing correction, volcano plots, PCA, and result evaluation, as well as reference outputs and an automated evaluator.
- Strengthened evaluator reliability through independent verification and abnormal-output detection mechanisms, reducing the potential for reward hacking.
- Designed a long-horizon Lean theorem-proving task requiring agents to perform strategy planning, theorem decomposition, step-by-step proof construction, compiler-guided error correction, and final proof assembly.
- Used the Lean kernel for deterministic machine verification of generated proofs, transforming real-world bioinformatics and formal mathematics workflows into reproducible and automatically scored agent evaluation tasks.

TECHNICAL SKILLS

Programming: Python, C, R, SQL, MATLAB, Neo4j, Linux

Machine Learning: PyTorch, TensorFlow, scikit-learn

LLM & Agent Systems: RAG, LangChain, Prompt Engineering, LLM Tool Use, MCP, Multi-Agent Systems orchestration, LLM routing