

# Ming Ma

Email: mam2022@ion.ac.cn — Phone: +86 13854225898

## Academic Qualifications

---

**University of Chinese Academy of Sciences** - Center for Excellence in Brain Science and Intelligence Technology

*Ph.D. Candidate in Computational Neuroscience* 2022.09–present

- Supervisor: Tianming Yang
- GPA: 3.72/4.0 — Honor: Graduate Merit Student
- Research interests: Computational Neuroscience, Interpretability of LLMs

**Shandong University** - School of Mechanical, Electrical and Information Engineering

*B.Eng. in Automation* 2019.03–2022.06

- GPA: 92.89/100 — Rank: 2/75
- Honors: First-Class Academic Scholarship, Honors Degree

**Beijing Institute of Technology** - Ruixin Academy of Classic Learning

*Exchange Student in Automation* 2019.09–2020.06

- GPA: 90.2/100 — Honor: First-Class Second-Campus Scholarship

**Naval Aviation University**

*Mechatronic Engineering* 2017.08–2019.03

- Rank: 3rd in battalion

## Research Experience & Publications

---

**DoVer: Intervention-Driven Auto-Debugging for LLM Multi-Agent Systems**

*Microsoft Research Asia* 2025

- First author: Proposed **DoVer**, upgrading passive error logging to an active *intervene–verify* loop. Built the supporting infrastructure (e.g., **AGDebugger**) with checkpointing on **Magentic-One** and **AG2**, enabling fully automated root-cause analysis and fix injection.
- Designed a multi-step intervention strategy that segments agent trajectories to isolate critical errors. This allows the system to progressively verify fixes and advance tasks through complex failure scenarios.
- On **GAIA** and **AssistantBench**, converted **18–28%** of failed cases into successes, establishing a general methodology for improving the robustness and debugging efficiency of multi-agent systems.

**A Hybrid Early-Exit Algorithm for Large Language Models Based on Space Alignment Decoding (SPADE)** 2025

- Co-first author: Proposed **SPADE** to solve "space drift" by aligning internal model states directly to the output. This enables the model to predict tokens from early layers without additional training, showing strong generalization.
- Introduced **SAD-Linear** with an adaptive early-exit policy that dynamically decides when to stop calculation. This reduces average inference depth by **55–70%** on benchmarks like **ARC** and **BoolQ** with negligible accuracy loss.
- Optimized the decoding process by compressing context into a minimal sequence, establishing a highly efficient, training-free acceleration paradigm for LLM deployment.

## Label Words as Local Task Vectors in In-Context Learning 2024

- Co-first author: Showed that in *out-of-distribution* and randomized ICL **abstract classification** tasks, LLMs rely on distributed **local task vectors** induced by label words rather than a single global task encoding.
- Validated the hypothesis via activation transfer/patching experiments; adapted *demixed PCA (dPCA)* from neuroscience to disentangle representational axes and performed vector-level ablations to establish **causal** effects.
- Characterized distinct behaviors in shallow vs. deep layers, offering an interdisciplinary lens for understanding ICL mechanisms.

## From End-to-end to Step-by-step: Learning to Abstract via Abductive Reinforcement Learning IJCAI 2025 (CCF A)

- Contributing author (in collaboration with Nanjing Univ. LAMDA Lab): Introduced a neuroscience-inspired *Drift Diffusion Model (DDM)* signal to detect *deadlock vs. exploration* states and guide switching decisions in the proposed **A2RL** framework.
- Designed and executed distributed delayed-reward experiments based on *RUD-DER*, along with ablations and comparisons; in sparse-reward, long-horizon **Grid-World** settings, validated that **A2RL** achieves stronger stability and robustness than conventional end-to-end RL.

## Intern Experiences

---

### Research Intern (LLM Agents)

*Microsoft Research Asia (MSRA)*

2025.07–2025.10

- **DoVer: Intervention-Driven Auto-Debugging for LLM Multi-Agent Systems:**  
See Research Experience & Publications section for details.
- **A Dual-Graph Memory Enhanced Deep Research Agent:**  
Built a *Co-evolving Dual-Graph* agent to improve long-horizon research tasks by separating report structure from factual knowledge. The system automatically identifies logical gaps to drive dynamic planning and targeted search, significantly enhancing the depth and accuracy of generated reports.

### Research Intern (Foundation Model Pretraining)

*Ant Group — Ling (Bailing) Foundation Model Team*

2025.10–present

- **Knowledge-conflict resolution and quality optimization:**  
Developed a pipeline to clean and synthesize data by resolving factual conflicts (e.g., using Wiki contradictory labels). Optimized corpus mixing ratios for pre-training, which improved *Ling-Mini*'s accuracy on benchmarks like *SimpleQA* and *WikiContradict* by **3–12%**.
- **Training strategy for structured table data:**  
Proposed an adaptive format strategy: converting simple tables to Markdown for better semantics while keeping HTML for complex ones to preserve structure. This approach, alongside an improved evaluation workflow, boosted the model's table reasoning capabilities by **3%**.
- **Evaluation pipeline bug fixes and standard alignment:**  
Fixed logical errors in perplexity (PPL) computation by upgrading to full-option evaluation. Resolved tokenization inconsistencies (e.g., handling leading spaces) to eliminate systematic bias, establishing a more rigorous and reproducible benchmark for model iteration.

## Projects

---

### **Computational Neuroscience Modeling: Decision-Process Simulation from Electrophysiological Signals**

*Neuromatch Academy - Computational Neuroscience*

2023.07

- **Model design & training:** Leveraged a public non-human-primate electrophysiology dataset to build and train a recurrent neural network (RNN) that simulates and predicts behavioural choices in a decision-making task.
- **Cross-modal representational analysis:** Employed representational similarity analysis (RSA) to align RNN hidden-state dynamics with recorded neuronal spike trains, revealing convergent decision-coding patterns shared by artificial networks and biological brains.

### **Bio-explainable Learning: Biologically Plausible CNN Training and Brain-Signal Analysis**

*Neuromatch Academy - NeuroAI*

2024.07

- **Biologically plausible optimization:** Replaced back-propagation with Hebbian local learning rules to train a convolutional neural network (Hebbian-CNN) for EEG classification, grounding the training process in neurobiological principles.
- **Model-brain alignment analysis:** Compared Hebbian-CNN, a standard CNN, and raw EEG signals in terms of classification accuracy and internal representations, demonstrating the representational gains afforded by biologically inspired learning.