Project 16
Weekly Update: Jul 6, 2026

Cards on the Table:

The Winning Move is Not in the Rules

♠ ♥ ♦ ♣

REU Student: Jack Lamberg

Sabbir Mollah, Sarinda Samarasinghe, Dr. Mubarak Shah

Problem Formulation



Real card-game agents should improve after every game.

1. Card-game skill depends on learning what cards, effects, and combos exist.
2. Rules define legal moves, but card effects have higher precedence over rules.
3. A move that looks good now may be bad with knowledge on what cards exist.

Our contribution:

1. LLM/VLM Understanding: Evaluate LLM/VLM performance on in-game questions.
2. LLM/VLM Gameplay: Evaluate LLM/VLM performance in playing full games.
3. Memory-based agent: Evaluate impact of accumulated memory on these metrics.

Memory Papers

- MemP: Exploring Agent Procedural Memory, (ACL Findings 2026)
- Meta-Policy Reflexion, (arXiv 2025)
- Coarse-to-Fine Grounded Memory for LLM Agent Planning, (EMNLP 2025)
- A-MEM: Agentic Memory for LLM Agents, (NeurIPS 2025)

Runnan Fang[♠♡*], Yuan Liang[♠*], Xiaobin Wang[♡], Jialong Wu[♡],
Shuofei Qiao[♠♡], Pengjun Xie[♡], Fei Huang[♡], Huajun Chen[♠], Ningyu Zhang[♠♣†]
[♠]Zhejiang University [♡]Alibaba Group
[♣]State Key Lab. for Novel Software Technology,
Nanjing University, P.R. China

MemP (ACL Findings 2026)

Summary:

- Concerned with giving LLM agents procedural memory
- Turning trajectories into reusable instructions
- Building, retrieving and updating memory

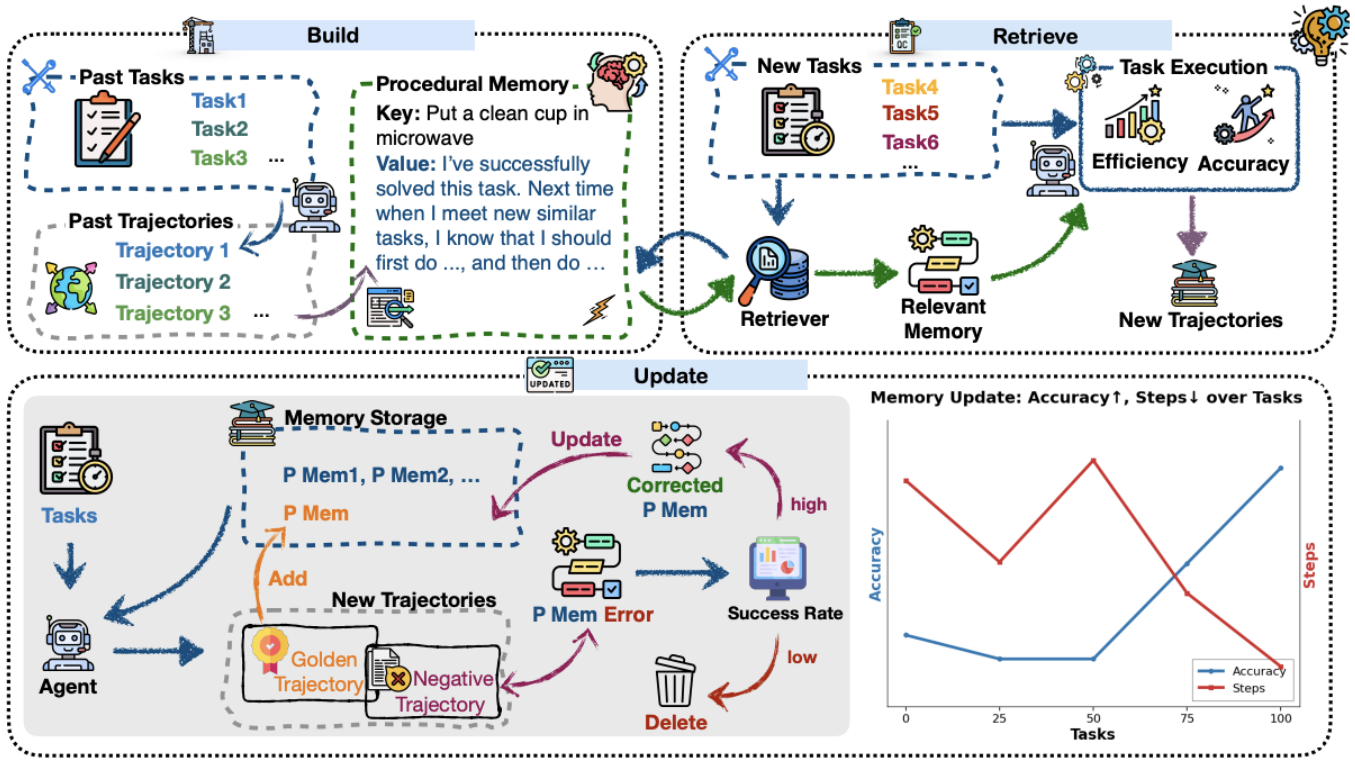

Sikuan Yan[*1,2], Xiufeng Yang[*3], Zuchao Huang[1], Ercong Nie[1,2],
Zifeng Ding[2,4], Zonggen Li[5], Xiaowen Ma[1], Jinhe Bi[1], Kristian Kersting[6],
Jeff Z. Pan[7], Hinrich Schuetze[1,2], Volker Tresp[1,2], Yunpu Ma[†1,2]
[1]Ludwig Maximilian University of Munich, [2]Munich Center for Machine Learning,
[3]Technical University of Munich, [4]University of Cambridge, [5]University of Hong Kong,
[6]Technical University of Darmstadt, [7]University of Edinburgh

Memory-R1 (ACL 2026)

Summary:

- Memory management using reinforcement learning
- Decide which memory to

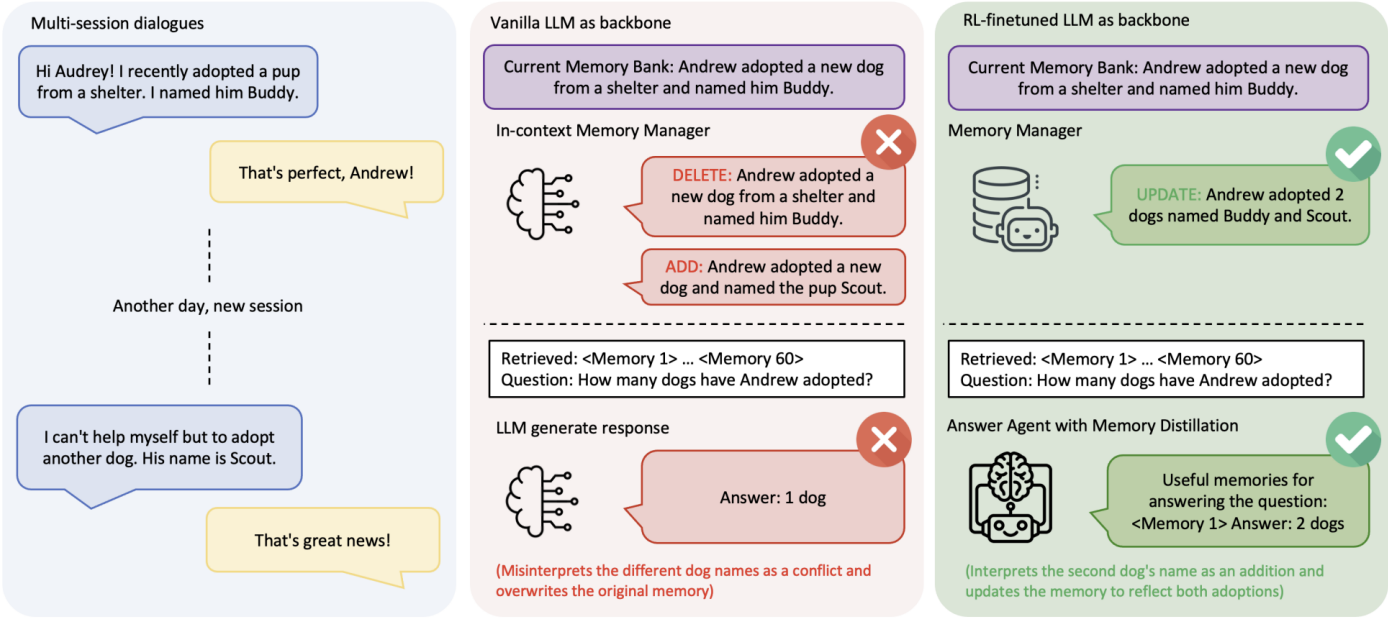

Wujiang Xu[1], **Zujie Liang**[2], **Kai Mei**[1], **Hang Gao**[1], **Juntao Tan**[1], **Yongfeng Zhang**[1,3]
[1]Rutgers University [2]Independent Researcher [3]AIOS Foundation
wujiang.xu@rutgers.edu

A-MEM (NeurIPS 2025)

Summary:

- Organize memory into a memory network
- Comprehensive memory notes
  - Description
  - Keywords
  - Tags
  - Links

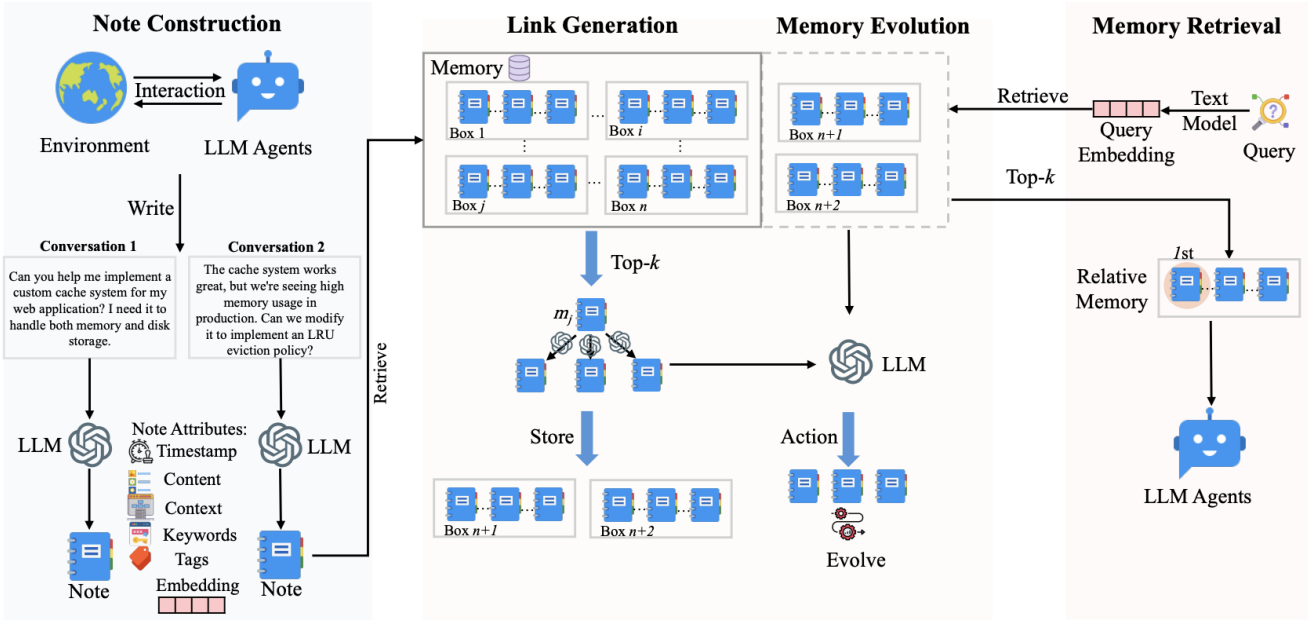

Proposed Memory Architecture

- The previous three methods can be applied to Card Games
  1. **Memp (Strategy Memory)**:
     - Converts game logs into reusable strategy procedures.
  2. **Memory-R1 (Memory Management):**
     - Decide whether to ADD, UPDATE, DELETE, or ignore each candidate memory.
  3. **A-MEM (Linking Memory Notes)**:
     - Stores memories as linked strategy notes with tags, keywords, and relationships.

Proposed Memory Note

- Each memory could have:
    1. Trigger conditions
    2. Recommended action pattern
    3. Exceptions
    4. Opponent archetype
    5. Evidence from past games
    6. Tags and links to related memories

Memory: Attacking under lethal pressure

1. Trigger conditions:
- opponent_next_turn_damage >= player_health
- player_has_attacker = true

2. Recommended action:
- Attack before slow value/resource plays.

3. Exceptions:
- Another move is better than attacking, gaining more tempo

4. Opponent Archetype:
- This move is good against an aggressive player

5. Evidence:
- +4 wins when followed
- -3 losses when ignored

6. Tags:
- #lethal-check
- #attack-timing
- #anti-aggro

Discussion on Pre-Built Bots

- Humans are better than the developed bots
  - Sarinda vs Random: 100% win-rate
  - Sarinda vs Rule-Based: 100% win-rate
  - Sarinda vs Greedy: 100% win-rate
  - **Sarinda vs MCTS-like: 100% win-rate**
- BUT, bots are currently better than the LLMs/VLMs
- **Result:** The bots are a useful intermediate evaluation to see if the models can play the game
  - If a model beats a bot more often, it would also improve against a human player

Models vs Bots (48 Games per matchup)

Win Percentage						
		Random	Rule-Based	Greedy	MCTS-like	Overall*
With Game History	QwenLLM-32B	83.3%	10.4%	16.7%	4.2%	28.65%
	Llama-8B	29.2%	4.2%	8.3%	0.0%	10.43%
	LLaVA-7B	91.7%	15.2%	27.1%	19.6%	38.40%
	InternVL-38B	85.4%	14.6%	8.3%	8.5%	29.20%
	QwenVLM-32B	91.5%	10.9%	15.9%	8.7%	31.75%

Week 7-8 Deliverables

Task	Week 6-7
Evaluate model vs MCTS	✅
Generate larger dataset for game understanding questions	✅
Complete evaluation on larger dataset	🟡
Play language models against each other	🟡
Literature review on memory	✅
Implement proposed memory module	🟡

Thanks for listening!

Questions?

