Generative AI and Application

How GenAI relates to VPPs

RMIT UNIVERSITY Ready for what's next

Generative AI and Application：How GenAI relates to VPPs

Guest speaker: A/Prof. Wei Peng

School of Engineering
RMIT University

Overview

- My Background
- Introduction to AI/GenAI
- GenAI Application and Challenges
- Related GenAI Research: LLM (Explainable AI) and Agentic AI (Workflow, Trajectory Learning, Multi-agent CDM)
- GenAI Application to VPPs

My Background

- 2025 – now: A/Professor, Principal Research Fellow in Industrial AI, School of Engineering, RMIT University
- 2019 – 2025: Research Director ITRC/AARC, AI Expert, BPIT, Huawei Technologies
- 2016 – 2019: Senior Lecturer, La Trobe University
- 2010 – 2016: Senior Data Scientist, AUSTRAC, Lenovo, Telstra
- 2008 – 2010: Research Fellow, RMIT University
- 2007 – 2008: Research Scientist, CSIRO ICT
- 2007: PhD from the University of Sydney in Agents in Design

Research Interests: *Large Language Models, Knowledge Representation, Cognitive Agents, and Embodied AI, along with their applications, 70+ papers, multiple patents*

Career Highlights: *research, innovation, and commercialization across both industry and academia*

Generative AI basics

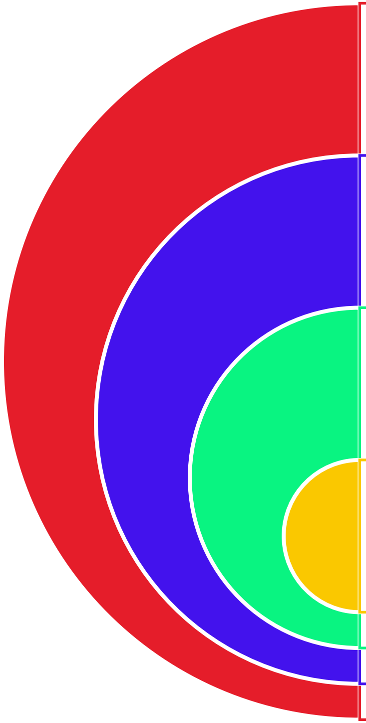

Artificial Intelligence	• Development of systems emulating human cognition in reasoning, learning, problem-solving and decision-making
Machine Learning	• Subset of AI that enables systems to learning from data and improve task performance without programming
Deep Learning	• Subset of ML that uses ANN with many layers (deep) to learn feature representations from large datasets to solve complex tasks
Generative AI	• A class of AI techniques generating new/original contents based on learned representations from huge datasets, often leveraging advanced neuron architecture like LLMs, VAEs, GAN, Diffusion models, etc.

AI Journey

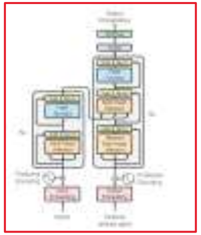

1st Winter

2nd Winter

1950-1972

Inception

1950 Turing Test
1956 Dartmouth Proposal
1958 First Perceptron, Lisp
1965 Eliza Chatbot, Fuzzy Sets
1968 First Knowledge-based System (Symbolic reasoning)
1969 First IJCAI at Stanford
1972 Prolog

1980-1987

Expert System

1980 First AAAI at Stanford
1982 Hopfield Network
1983 Soar Cognitive Arch.
1986 Backpropagation for MLPs

1991-2017

Machine Learning

1991 DART in 1st Golf War
1993 Behavior-based Robots, Agents
1994 Soft Computing (Fuzzy + ANN, GE)
1997 Deep Blue (IBM) defeat Kasparov, LSTM, 1st RoboCup
2002 iRobot's Roomba
2004 NASA Spirit and Opportunity on Mars
2011 Apple Siri (iPhone 4s)
2012 AlexNet
2016 AlphaGo defeated Lee Sedol

2017-2025

GenAI: DL/LLM

2017 Google Transformer, AlphaZero, OpenAI Bot Win Dota
2020 AlphaFold 2 wins CASP, OpenAI GPT3
2021 OpenAI DALL-E
2022-23 OpenAI ChatGPT, GPT4
2024 OpenAI Sora, Nobel Prize (AlphaFold)
2025 DeepSeek, Agentic

2025+

AGI/ASI

Agentic
Artificial General AI
Artificial Super AI
Singularity

Hype Cycle for GenAI in 2024 (to watch out)

➢ 80% enterprises will deploy GenAI in production in 2026 (Gartner, 2024)

Strengths: automation, reasoning, creativity

Challenges: hallucination, factuality, explainability

Four Main Core Technologies:

1. GenAI models: Bigger models, Embedding, Domain-specific, Edge GenAI, AGI

2. GenAI engineering: AI Trism, Disinformation, Orchestration, GraphRAG

3. GenAI application and use cases: Virtual assistant, GenAI software engineering, Autonomous agent, Synthetic data

4. GenAI enablement: Workload accelerators, AI simulation, Supercomputing, Self-supervised, Transfer learning

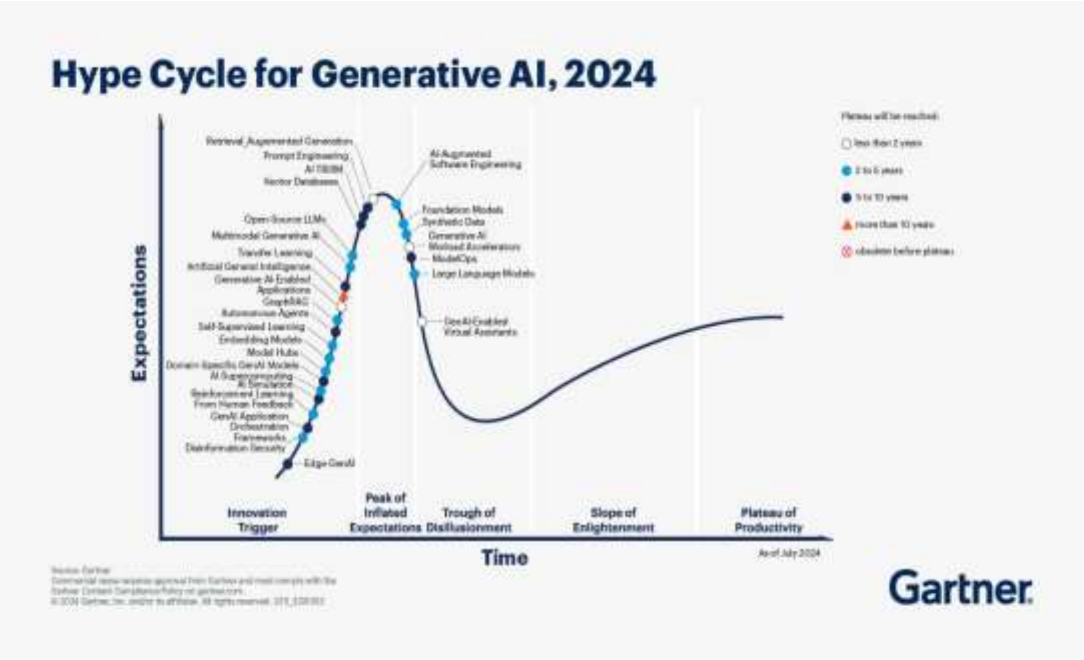

https://www.gartner.com/en/articles/hype-cycle-for-genai

Value Proposition of GenAI

The potential impact of generative AI can be evaluated through two lenses.

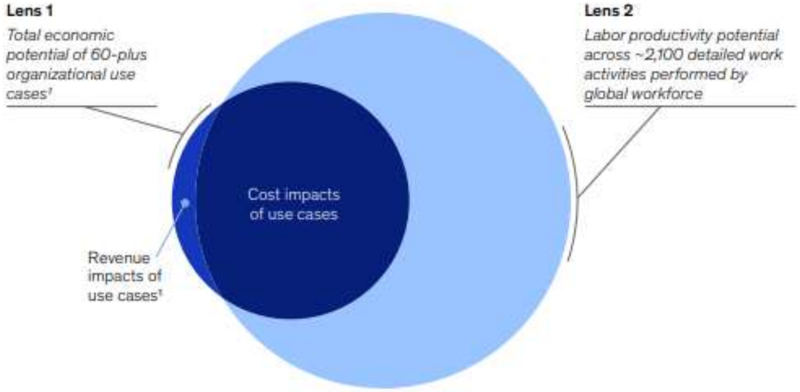

[1]For quantitative analysis, revenue impacts were recast as productivity increases on the corresponding spend in order to maintain comparability with cost impacts and not to assume additional growth in any particular market.

Generative AI could create additional value potential above what could be unlocked by other AI and analytics.

AI's potential impact on the global economy, $ trillion

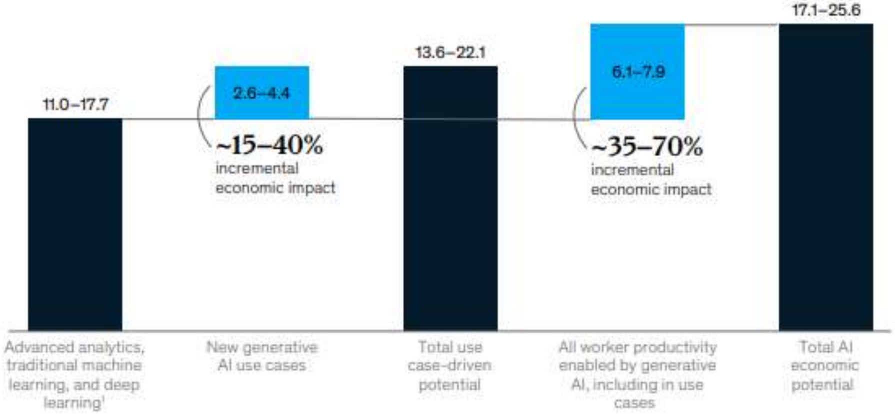

McKinsey & Company, 2023

Generative AI Applications in Enterprise

➢ **Automotive:** Toyota uses GenAI to produce initial vehicle sketches, reduce time-to-design, minimize unexpected late design changes

➢ **Healthcare:** Mayo Clinic leverages Medical Chatbots to conduct dynamic interviews with patients and provide personalized care recomm.

➢ **Banking:** Ally's contact center assistant to automate note-taking and summarizing customer calls, reduce manual call services, focus associates On customer interactions

➢ **Benefits:** Increase revenue, efficiency and othe nonfinancial value, manage risk

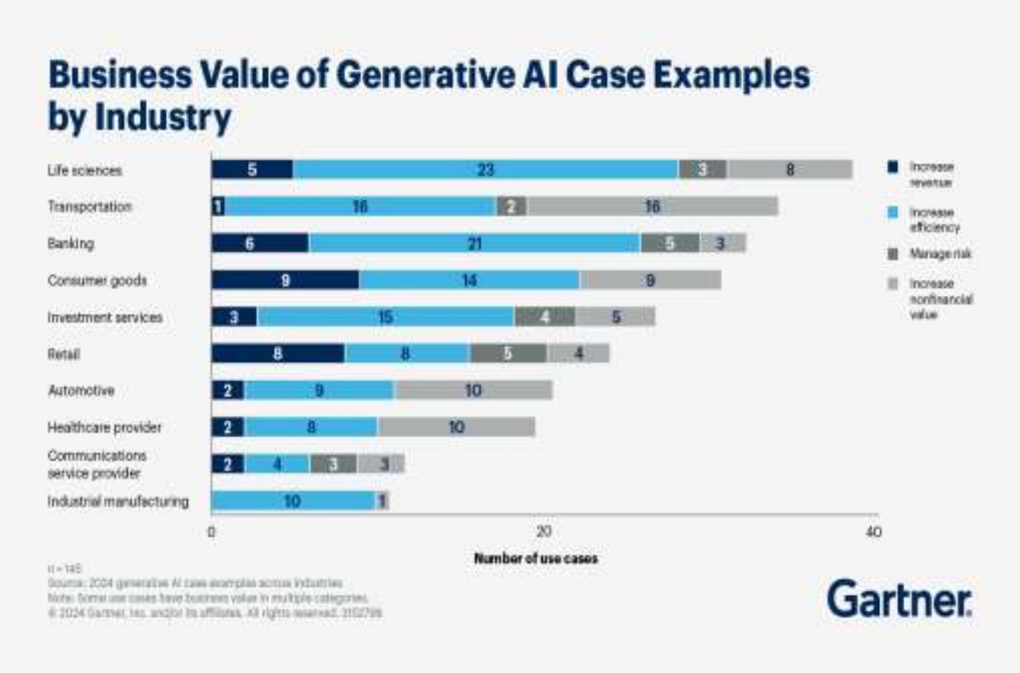

https://www.gartner.com/en/articles/generative-ai-use-cases

Low-hanging Fruit of GenAI Application

Using generative AI in just a few functions could drive most of the technology's impact across potential corporate use cases.

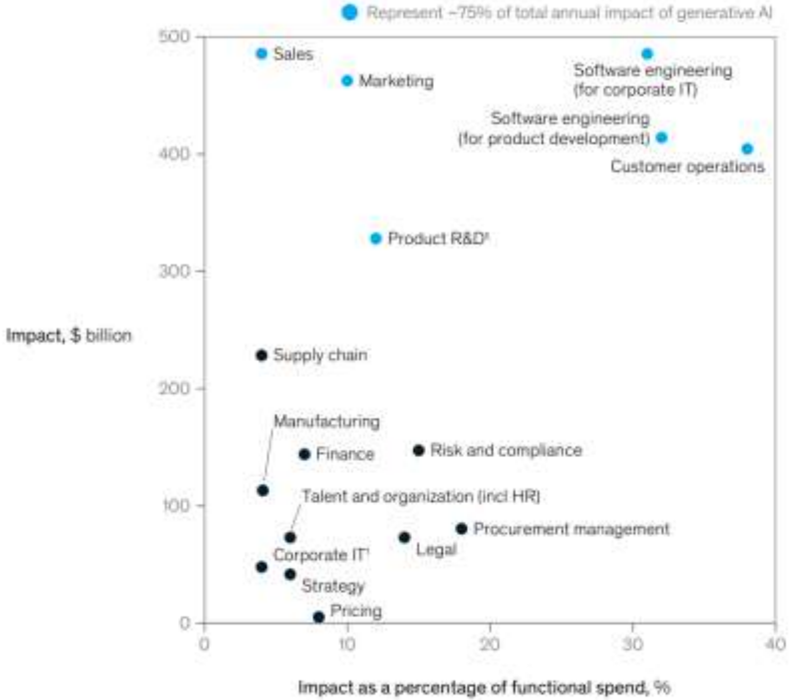

McKinsey & Company, 2023

Generative AI use cases will have different impacts on business functions across industries.

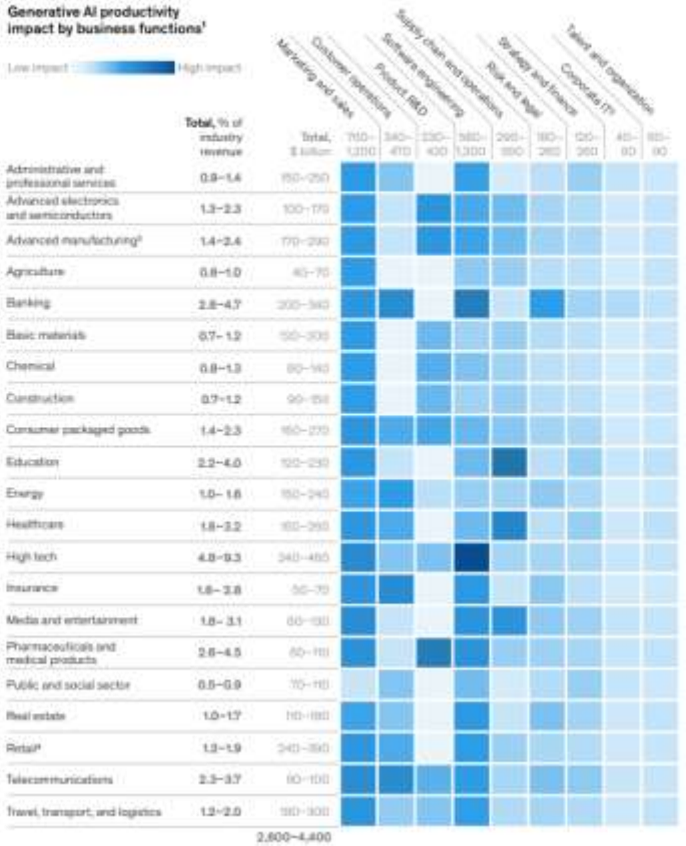

Low-hanging Fruit of GenAI Application (cont.)

1. Customer Service & Support: chatbots and virtual assistant, multilingual NMT, autogen contents for email and tickets
2. Marketing and Content Generation: Ads, Email, product manual
3. Domain (internal) Knowledge QA: Gen-AI assistant trained on internal docs (RAG/KAG), HR on-boarding, training, policy, doc summarization, helpdesk
4. Code Generation and Software Dev: Code generation, automated documentation, unit test, de-bug …
5. Meeting and Communications: voice-to-text assistants (ASR), meeting transcripts and summarization, action items
6. Workflows: LLM-based Agentic AI for tasks automation enabling above-mentioned key functions

Challenges of GenAI

- **Hallucination and Accuracy**

GenAI produces hallucinations due to limitation of data, training, contextual influence. It is the nature of GenAI. Invest in XAI technologies to detect and rectify hallucination, enhance interpretability of innate mechanism. GenAI fits in training data will have to deal with generalization across unseen real-world problems.

- **High Computation Cost**

SFT on LLMs very costly in computational cost, need to seek more efficient ways in post-training, invest in model compression, distillation or inference time computing (in-context learning) or activation steering. Adaptive learning merges.

- **Domain Data and Knowledge Gap**

LLM pre-trained on human general knowledge. Gap in domain knowledge due to lack of data. Need to ensure data quality and availability, Leverage domain data generation and knowledge integration, like domain data enhancement, RAG/KAG, Agentic AI

- **AiTRISM**

Tackling trust, risk, security of GenAI. Protect sensitive project and ensure the security of AI, ensure ethical and responsible AI by protecting data privacy

My Research Stream I: XAI, Low-cost Intervention

- **Detecting Hallucination:**

Assessing Factual Reliability of Large Language Model Knowledge, in *Proceedings of NAACL 2024, Oral Presentation*, Association for Computational Linguistics

A survey on hallucination in large vision-language models, 2024, arXiv preprint arXiv:2402.00253

- **Inconsistency (Semantics, Preferential Ranks):**

Enhancing Semantic Consistency of Large Language Models through Model Editing: An Interpretability-Oriented Approach, In *Findings of the Association for Computational Linguistics: ACL 2024,* Association for Computational Linguistics

Measuring the Inconsistency of Large Language Models in Preferential Ranking, In *Proceedings of the 1st Workshop on Towards Knowledgeable Language Models (KnowLLM 2024),* Association for Computational Linguistics

- **Activation Steering:**

Semantics-Adaptive Activation Intervention for LLMs via Dynamic Steering Vectors, In Proceedings of ICLR 2025, https://arxiv.org/abs/2410.12299

LF-Steering: Latent Feature Activation Steering for Enhancing Semantic Consistency in Large Language Models, 2025, arXiv preprint arXiv:2501.11036

Bifurcated Pathway to AGI (Data-driven)

1. Human knowledge build (Cybernetics, KBS, to Data and Computing Power, LLM/scale law/transient learning on features for tasks)

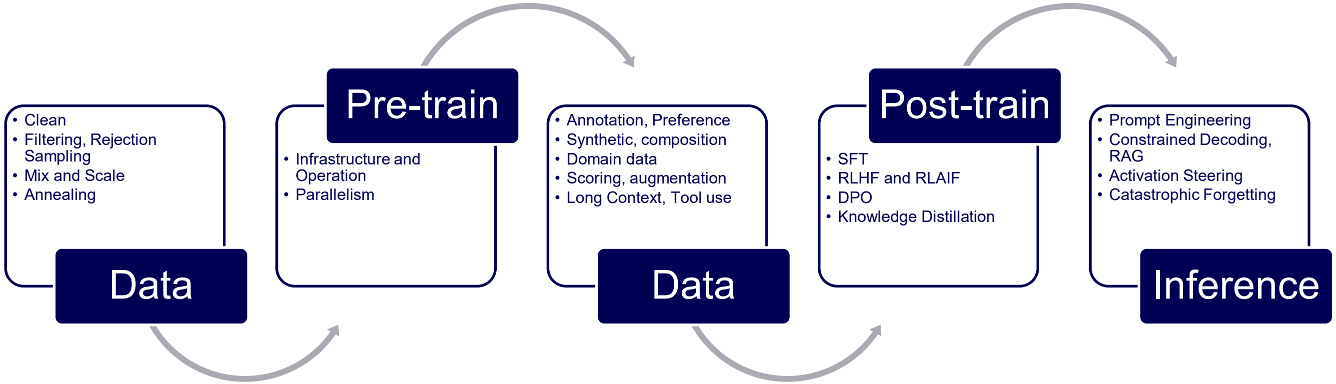

We try to scale model in development to host all human knowledge and capability by feeding mega data …

Bifurcated Pathway to AGI (Experience-driven)

2. Continual learning from experience

Rich Sutton's new path for AI : "… RL in AI, we don't have methods to learn continuously except for the linear case …" https://www.youtube.com/watch?v=NvfK1TkXmOQ

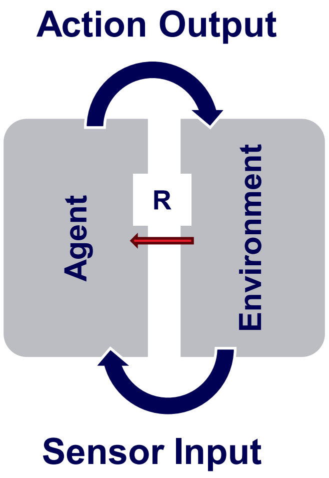

DRL Agent

Learn from Interactions Challenges:

- Generalization: beyond the env. trained
- Inefficiency: learn from limited examples
- Catastrophic Forgetting: learn without forgetting priori knowledge
- Reward Misalignment: multi-objective, human values
- Lack World Model and Abstraction: reason on concept

Learning from Interactions (Grounding Problem)

- **Grounding:** intrinsic process of assigning meanings to symbols/words/vectors/concepts by referencing to real world experience (objects, events).
- **Symbolic grounding** (Harnad, 1990): how can the semantic interpretation of a formal symbol system be made intrinsic to the system, rather than just parasitic on the meaning in our heads?
- **Representation grounding** (Chalmers, 1992): how can a representation in a computational system possess true meaning?
- **Concept grounding** (Dorffner & Prem, 1993): design a cognitive model (connectionism) only interfacing with its environment using sensor and motor signals; any concept of the system develops through self-organization based on adaptive interaction with the environment (besides given meta-level representation like innate architecture) – is grounded in Harnard's sense.

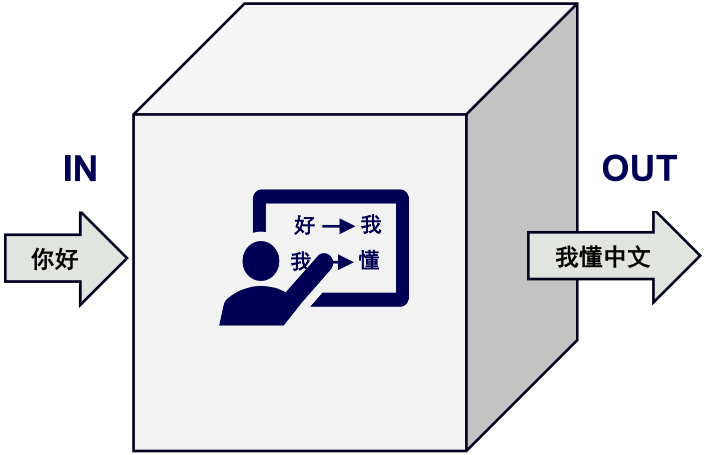

Searle's Chinese Room

Learning from Interactions (Situated Intelligence)

- Rodney Brooks' Intelligence without Representation (Brooks, 1991): no traditional representation, intelligence from sensor motor interaction with the environment, behavior-based model

Take Aways: Learn from situated sensor motor coordination to generate complex behaviors

Challenges: Limited memory, planning, reasoning capability, simplistic world model

Key Insight: ground disembodied intelligence (i.e., LLM) in interactions to develop concepts (levels of abstractions) in self-organized manner

Agents Vs Agentic AI

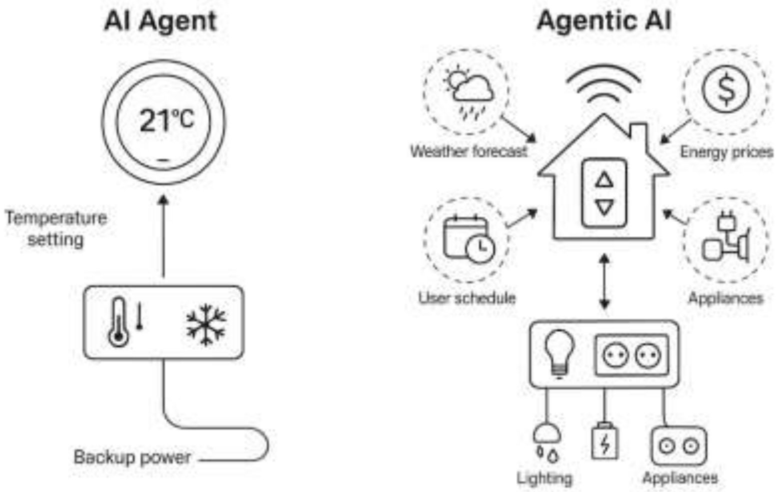

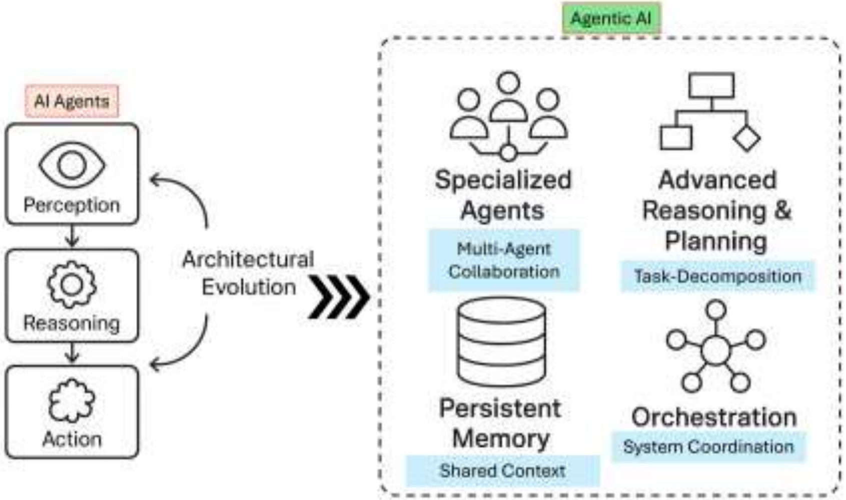

Left: A single-task AI Agent.
Right: A multi-agent, collaborative Agentic AI system

Left: domain-specific prompts, context-aware (ReACT)
Right: Shift from isolated perception-reasoning-action loops to collaborative and self-evaluative multi-agent workflows enables agents to reflect, learn and improve over time

AI Agents vs. Agentic AI (Sapkota et al. 2025)

My Research Stream II: Agentic AI (Workflow Agents)

➢ **Workflow Agents (NL2Workflow):**

WorkTeam: Constructing Workflows from Natural Language with Multi-Agents, In *Proceedings of NAACL 2025, Industry Track, pages 20-35*, https://aclanthology.org/2025.naacl-industry.3/

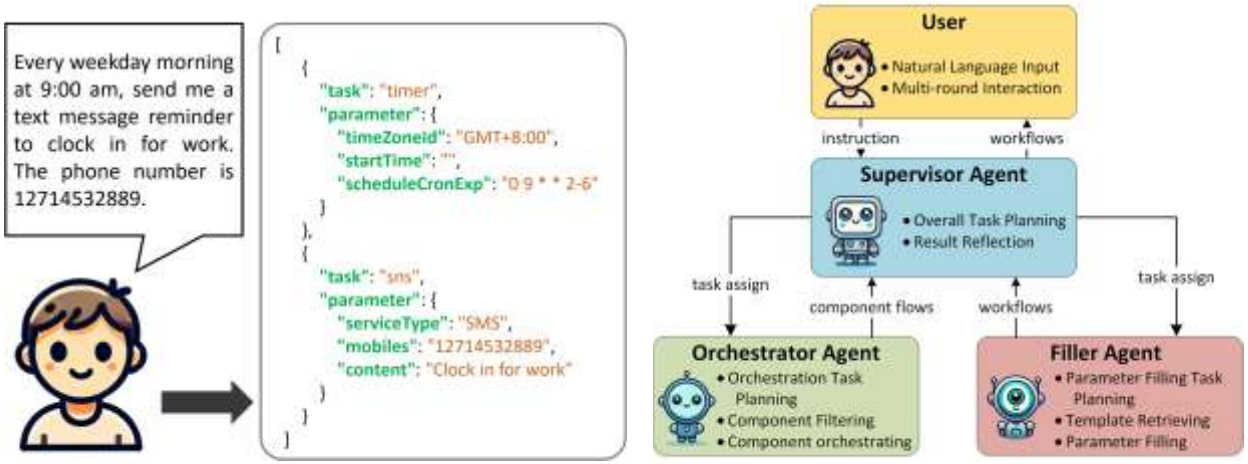

Methods	EMR (%)	AA (%)	PA (%)
GPT-4o	18.1	71.4	56.3
Qwen2.5-72B-Instruct	12.7	66.9	51.5
Qwen2.5-7B-Instruct	3.5	25.4	19.9
LLaMA3-8B-Instruct	1.6	19.4	16.6
RAG (Ayala and Bechard, 2024)	24.1	77.8	60.3
WorkTeam (ours)	**52.7**	**88.9**	**73.2**

WorkTeam deployed with a real-world benchmark of 93% accuracy

My Research Stream II: Agentic AI (Trajectory Learning)

➢ **Agent Learning from Interaction:**

AgentBank: Towards Generalized LLM Agents via Fine-Tuning on 50000+ Interaction Trajectories, *In Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 2124–2141, Association for Computational Linguistics

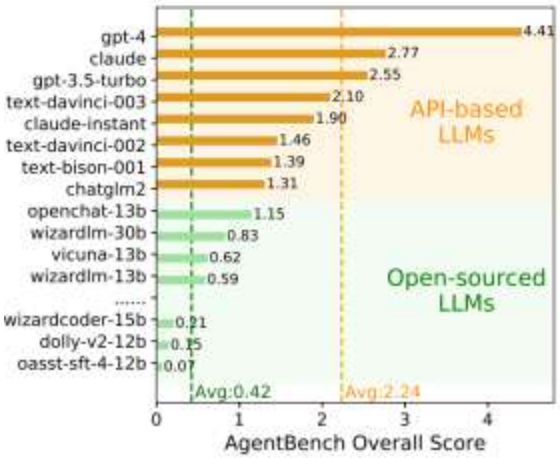

LLMs suffer even for SOTA commercial tools

Expert Trajectories

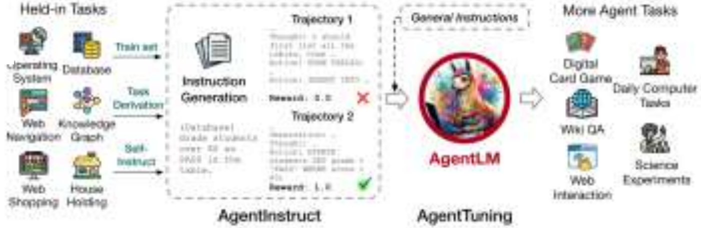

AgentTuning: only keep successful trajectories for training

Agents solve only simple tasks, hard to generalize

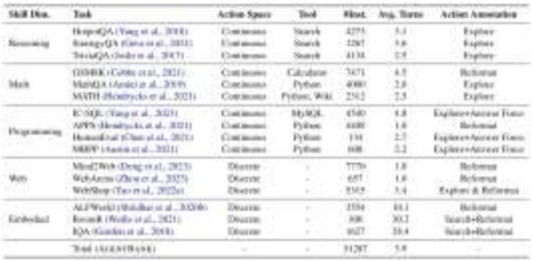

Skill Dim.	Task	Action Space	Tool	#Inst.	Avg. Turns	Action Annotation
Reasoning	HotpotQA (Yang et al., 2018)	Continuous	Search	4275	3.1	Explore
	StrategyQA (Geva et al., 2021)	Continuous	Search	2287	3.6	Explore
	TriviaQA (Joshi et al., 2017)	Continuous	Search	4138	2.5	Explore
Math	GSM8K (Cobbe et al., 2021)	Continuous	Calculator	7471	4.5	Reformat
	MathQA (Amini et al., 2019)	Continuous	Python	4080	2.0	Explore
	MATH (Hendrycks et al., 2021)	Continuous	Python, Wiki	2312	2.5	Explore
Programming	IC-SQL (Yang et al., 2023)	Continuous	MySQL	4540	4.8	Explore+Answer Force
	APPS (Hendrycks et al., 2021)	Continuous	Python	4400	1.0	Reformat
	HumanEval (Chen et al., 2021)	Continuous	Python	134	2.7	Explore+Answer Force
	MBPP (Austin et al., 2021)	Continuous	Python	608	2.2	Explore+Answer Force
Web	Mind2Web (Deng et al., 2023)	Discrete	-	7770	1.0	Reformat
	WebArena (Zhou et al., 2023)	Discrete	-	657	1.0	Reformat
	WebShop (Yao et al., 2022a)	Discrete	-	3393	3.4	Explore & Reformat
Embodied	ALFWorld (Shridhar et al., 2020b)	Discrete	-	3554	10.1	Reformat
	RoomR (Weihs et al., 2021)	Discrete	-	308	30.2	Search+Reformat
	IQA (Gordon et al., 2018)	Discrete	-	1627	18.4	Search+Reformat
	Total (AGENTBANK)	-	-	51287	5.8	-

We propose AgentBank to host 50000+ interaction trajectories

My Research Stream II: Agentic AI (Trajectory Learning cont.)

➢ **Agent Learning from Interaction:**

AgentBank: Towards Generalized LLM Agents via Fine-Tuning on 50000+ Interaction Trajectories, *In Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 2124–2141, Association for Computational Linguistics

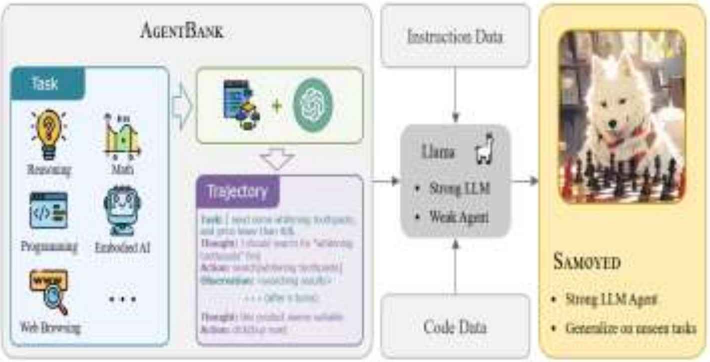

Model	Held-in Tasks						Held-out Tasks					
	Reason	Math	Program	Web	Embodied	Avg.	Reason	Math	Program	Web	Embodied	Avg.
Closed-Source Model												
GPT-4	61.6	73.0	54.9	40.6	77.8	59.8	41.6	51.0	69.4	69.4	36.4	53.6
GPT-3.5-Turbo	41.0	41.5	51.2	42.0	10.5	40.2	32.0	32.0	54.8	66.7	21.2	41.3
7B Open-Source Model												
Llama-2-7B-Chat	4.0	7.5	2.5	13.9	0.0	6.2	4.0	8.0	7.0	0.4	7.8	5.5
Vicuna-7B	29.0	2.0	19.0	24.2	6.0	17.1	8.8	14.0	19.0	18.2	12.8	14.6
CodeLlama-7B	3.5	3.5	1.5	24.8	0.0	7.4	1.0	13.0	21.8	**41.3**	5.5	10.5
AgentLM-7B	29.5	10.0	12.0	**37.2**	63.4	26.7	19.2	13.0	50.5	13.5	13.3	21.9
Agent-FLAN-7B	31.0	10.5	13.1	35.4	**65.3**	27.3	22.2	11.0	53.1	17.9	14.1	23.7
Samoyed-7B	**48.0**	**30.5**	**41.6**	35.4	61.2	**41.6**	**32.0**	**18.0**	**59.2**	24.2	**14.2**	**29.5**
13B Open-Source Model												
Llama-2-13B-Chat	12.5	10.5	8.2	11.2	0.0	9.4	9.6	11.0	33.0	17.6	7.3	15.7
Vicuna-13B	25.5	6.5	30.4	34.2	2.2	21.7	24.8	17.0	37.0	34.2	14.8	25.6
CodeLlama-13B	13.5	18.5	5.1	15.3	0.0	11.7	6.4	16.0	11.1	**46.5**	5.5	17.1
AgentLM-13B	38.0	15.5	22.8	38.1	52.2	30.8	20.8	13.0	46.6	21.6	14.6	23.5
Samoyed-13B	**54.5**	**38.5**	**55.4**	**40.9**	**72.4**	**50.1**	**35.0**	**23.0**	**62.4**	38.9	**18.4**	**35.5**

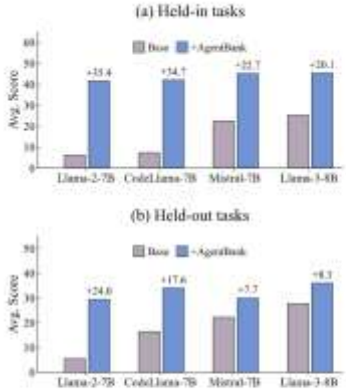

Organize trajectories into multi-turn dialogues, mix general domain instructions and codes, utilize failure trajectories and propose the exploration-based trajectory optimization (ETO) method to learn the task-solving process, leading to significant performance gains.

My Research Stream II: Agentic AI (Trajectory Learning cont.)

➢ **Agent Learning from Interaction:**

Watch Every Step! LLM Agent Learning via Iterative Step-Level Process Refinement, In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024),* pages 1556-1572, Association for Computational Linguistics

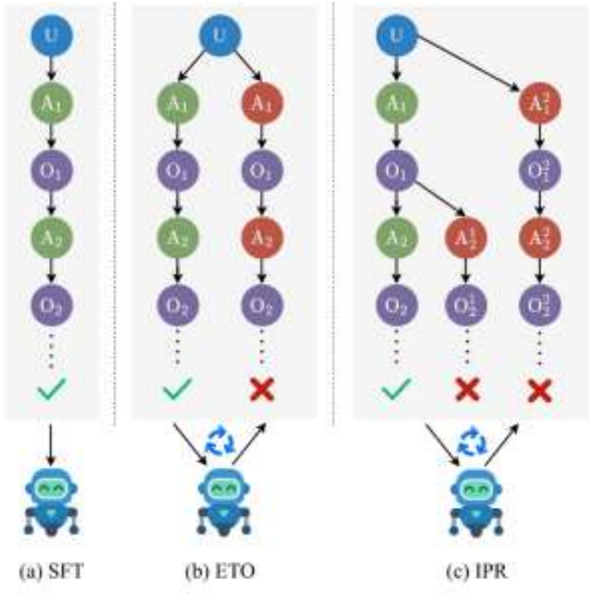

➢ **Agents start to learn from Interactions and explorations: from SFT on trajectories to ETO (SAMOYED)**

Treat an entire trajectory as single entity during training and prioritize the final reward of a trajectory over the process, thus overlooking exploitable information throughout interaction process.

We need to consider step level optimization

My Research Stream II: Agentic AI (Trajectory Learning cont.)

➢ Agent Learning from Interaction:

Watch Every Step! LLM Agent Learning via Iterative Step-Level Process Refinement, In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024),* pages 1556-1572, Association for Computational Linguistics

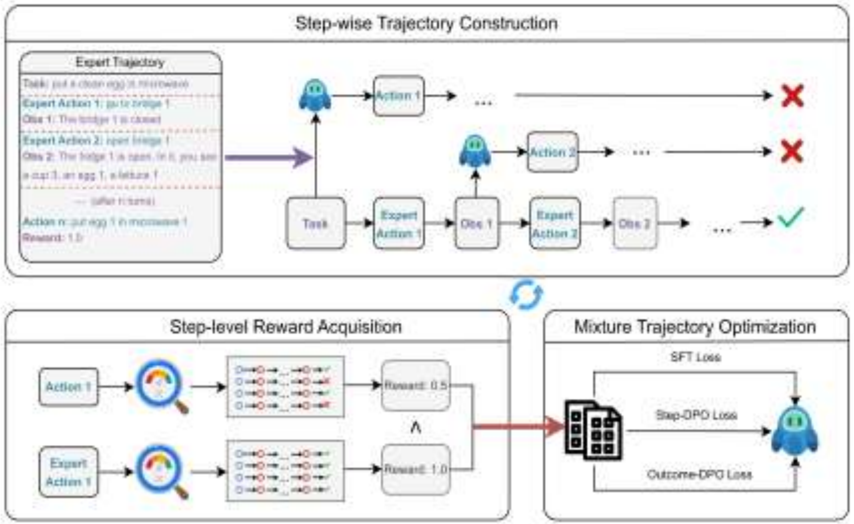

Design Step-level Process Refinement: Step-level Reward Acquisition and Iterative Agent Optimization.

Monte Carlo method to estimate rewards via sampling N trajectories to construct step award.

$$\{e^{(i)}|i=1,...,N\} = MC^{\pi_s}(e_{t-1};N),$$

$$r_s(s_t,a_t) = \begin{cases} \frac{1}{N}\sum_{i=1}^{N} r_o(u,e^{(i)}), & \text{for } t<n \\ r_o(u,e_n), & \text{for } t=n \end{cases}$$

$$\mathcal{L} = \mathcal{L}_{\text{o-DPO}} + \mathcal{L}_{\text{s-DPO}} + \mathcal{L}_{\text{SFT}}$$

- Outcome-DPO Loss

$$\mathcal{L}_{\text{o-DPO}} = -\mathbb{E}_{(u,e_n^w,e_m^l)\sim\mathcal{D}_o}\left[\log\sigma\left(\beta\log\frac{\pi_\theta(e_n^w|u)}{\pi_{ref}(e_n^w|u)} - \beta\log\frac{\pi_\theta(e_m^l|u)}{\pi_{ref}(e_m^l|u)}\right)\right],$$

- Step-DPO Loss

$$\mathcal{L}_{\text{s-DPO}} = -\mathbb{E}_{(e_{t-1},e_{t:n}^w,e_{t:m}^l)\sim\mathcal{D}_s}\left[\log\sigma\left(\beta\log\frac{\pi_\theta(e_{t:n}^w|e_{t-1})}{\pi_{ref}(e_{t:n}^w|e_{t-1})} - \beta\log\frac{\pi_\theta(e_{t:m}^l|e_{t-1})}{\pi_{ref}(e_{t:m}^l|e_{t-1})}\right)\right],$$

- Supervised Loss

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{(u,e_n^w,e_m^l)\sim\mathcal{D}_o}\left[\log\pi_\theta(e_n^w|u)\right],$$

My Research Stream II: Agentic AI (Trajectory Learning cont.)

➢ Agent Learning from Interaction:

Watch Every Step! LLM Agent Learning via Iterative Step-Level Process Refinement, In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024),* pages 1556-1572, Association for Computational Linguistics

Paradigm	Models	WebShop	InterCodeSQL	ALFWorld		Average
				Seen	Unseen	
Prompt-based	GPT-4 (Achiam et al., 2023)	63.2	38.5	42.9	38.1	45.7
	GPT-3.5-Turbo (Ouyang et al., 2022)	62.4	37.8	7.9	10.5	29.7
	Llama-2-7B (Touvron et al., 2023)	17.9	4.0	0.0	0.0	5.5
Outcome Refinement	Llama-2-7B + SFT (Chen et al., 2023)	60.2	54.9	60.0	67.2	60.6
	Llama-2-7B + PPO (Schulman et al., 2017)	64.2	52.4	22.1	29.1	42.0
	Llama-2-7B + RFT (Yuan et al., 2023)	63.6	56.3	62.9	66.4	62.3
	Llama-2-7B + ETO (Song et al., 2024)	67.4	57.2	68.6	72.4	66.4
Process Refinement	Llama-2-7B + Step-PPO	64.0	60.2	65.7	69.4	64.8
	Llama-2-7B + IPR (ours)	**71.3**	**61.3**	**70.3**	**74.7**	**69.4**

Training Scheme	WebShop	InterCodeSQL	ALFWorld
w/o o-DPO	70.2	59.3	72.4
w/o s-DPO	66.4	58.0	70.2
w/o SFT	61.8	31.7	64.9
Iteration=1	63.6	56.6	68.7
Iteration=2	63.7	58.2	70.2
Iteration=3	68.2	59.2	**74.7**
Iteration=4	**71.3**	**61.3**	73.5
Iteration=5	68.1	57.9	71.4

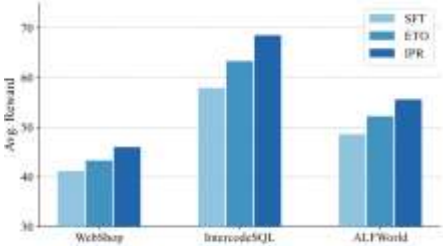

Conclusion：

- ➢ Agent learns from interaction via trajectory with step awards
- ➢ Learning from failure actions
- ➢ Automated process reward acquisition
- ➢ Step level process supervision via mixture trajectory optimization
- ➢ Enhanced performance on three benchmarks
- ➢ Generalizable on unseen hold out

Limitation：

- ➢ Overfitting with limited data (need to leverage AgentBank data)
- ➢ MC method constrained by sample size
- ➢ Consider GPT 4 to label process supervision data

My Research Stream II: Agentic AI (Multi-Agents CDM)

➢ **Multi-agents Collective Decision Making (CDM):**

An Electoral Approach to Diversify LLM-based Multi-Agent Collective Decision-Making, *in Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP 2024),* pages 2712–2727, Association for Computational Linguistics, https://aclanthology.org/2024.emnlp-main.158/

Table 4: Full list of 52 surveyed LLM-based multi-agent collaboration works.

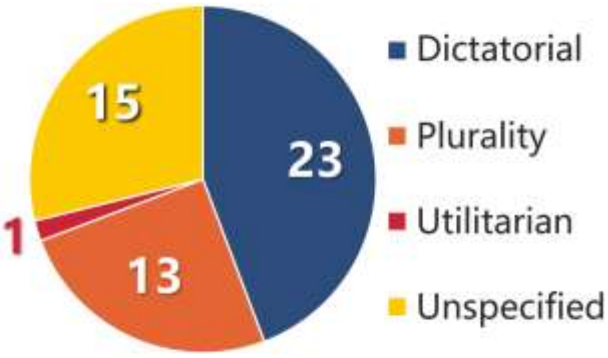

52 multi-agent collaboration frameworks: lack of diversity in Collective Decision-making (CDM)

CDM Method	Major -ity	Mono -tonic	Consis -tency	IIA	Cond -orcet	Ballot type
Dictatorial (Blind)	✗	✓	✓	✓	✗	Ranking
Range Voting	✗	✓	✓	✓	✗	Scores
Plurality	✓	✓	✓	✗	✗	Single*
Borda Count	✗	✓	✓	✗	✗	Ranking
IRV	✓	✗	✗	✗	✗	Ranking
Ranked Pairs	✓	✓	✗	✗	✓	Ranking

Table 1: Criteria compliance of some typical CDM methods. *Range Voting* can be viewed as a special *utilitarian* method. **IIA** denotes *Independence from Irrelevant Alternatives*. *Single ballots can be derived from ranking ones. Find some examples in Appendix D.

Kenneth Arrow's Social Choice Theory

Diversifying CDM in LLM MAS

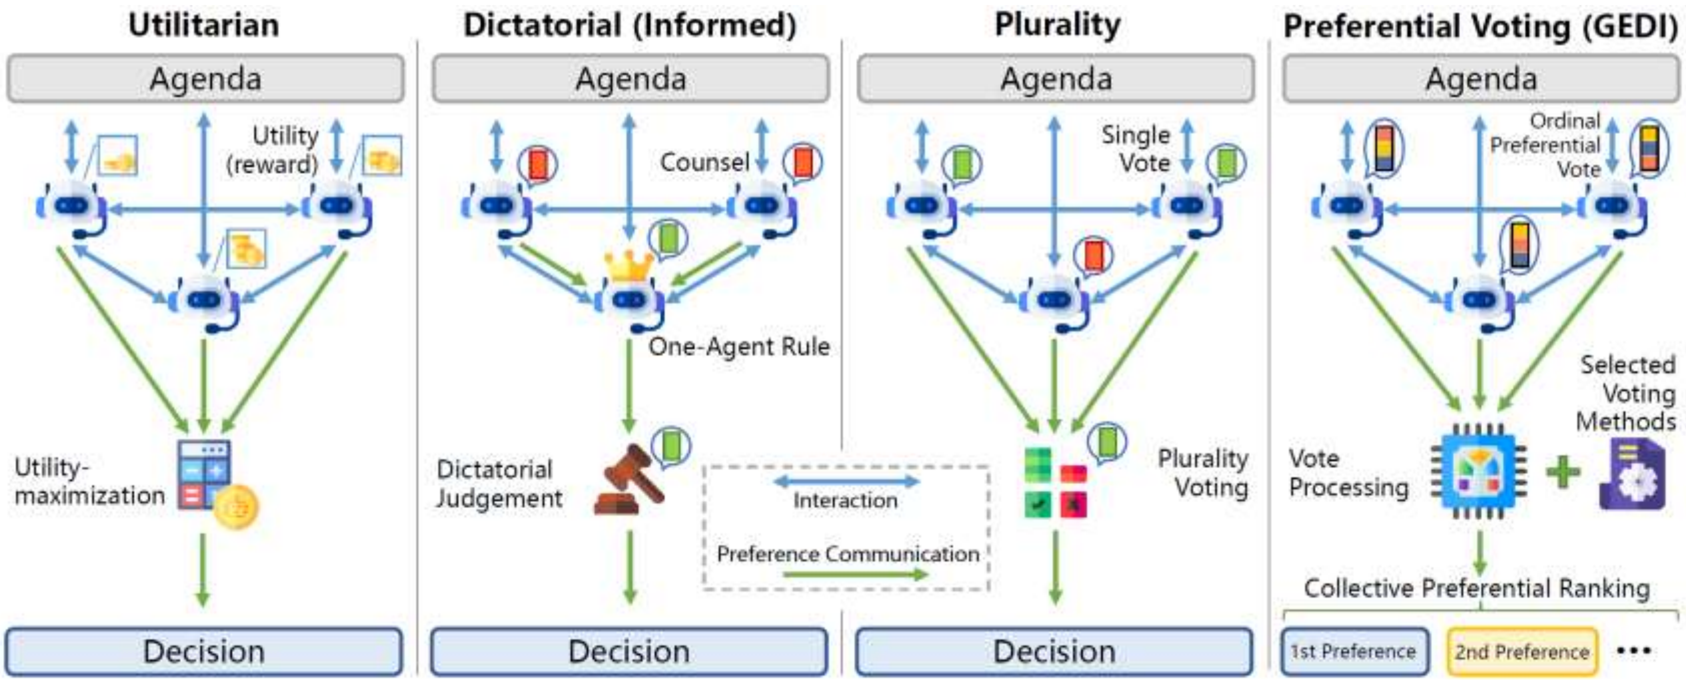

Key Findings

Base Model	Rand.	Score	Dictatorial-based			Ordinal Ranking					
MMLU	Rand.	Range Voting	Blind Dicta.	Informed Dicta.	Mis-Informed Dicta.	Plurality	Bucklin	Borda Count	IRV	Minimax	Ranked Pairs
mistral-7b	24.8	51.8 (-4.6)	56.4	55.9 (-0.5)	36.1 (-20.3)	56.8 (+0.4)	57.1 (+0.7)	56.9 (+0.5)	56.9 (+0.5)	57.0 (+0.6)	57.0 (+0.6)
llama-3-8b	25.0	37.7 (-7.3)	45.0	36.5 (-8.5)	32.2 (-12.8)	45.9 (+0.9)	46.4 (+1.4)	46.3 (+1.3)	45.7 (+0.7)	45.9 (+0.9)	46.0 (+1.0)
glm-4-9b	25.2	61.3 (-0.4)	61.7	54.3 (-7.4)	**53.0** (-8.7)	**64.6** (+2.9)	**64.5** (+2.8)	64.1 (+2.4)	**64.9** (+3.2)	**64.4** (+2.7)	**64.6** (+2.9)
llama-3-70b	25.3	74.9 (+1.6)	73.3	70.1 (-3.2)	62.6 (-10.7)	73.9 (+0.6)	73.8 (+0.5)	73.7 (+0.4)	73.9 (+0.6)	73.9 (+0.6)	73.9 (+0.6)
qwen-2-72b	25.1	69.2 (-0.5)	69.7	69.7 (±0.0)	39.5 (-30.2)	70.0 (+0.3)	69.9 (+0.2)	70.0 (+0.3)	69.9 (+0.2)	69.9 (+0.2)	69.9 (+0.2)
qwen-1.5-110b	25.0	71.3 (-1.5)	72.8	73.0 (+0.2)	46.3 (-26.5)	72.9 (+0.1)	72.9 (+0.1)	72.7 (-0.1)	72.9 (+0.1)	72.9 (+0.1)	72.9 (+0.1)
gpt-3.5	24.9	63.0 (+2.2)	60.8	**64.7** (+3.9)	36.9 (-23.9)	**65.9** (+5.1)	**65.5** (+4.7)	**65.6** (+4.8)	**65.6** (+4.8)	**65.6** (+4.8)	**65.6** (+4.8)
gpt-4	25.0	**80.7** (+5.1)	75.6	**82.1** (+6.5)	**70.9** (-4.7)	**82.5** (+6.9)	**81.9** (+6.3)	**81.9** (+6.3)	**81.9** (+6.3)	**81.9** (+6.3)	**81.9** (+6.3)
MMLU-Pro											
mistral-7b	9.6	20.9 (-9.0)	29.9	27.7 (-2.2)	15.6 (-14.3)	31.7 (+1.8)	30.7 (+0.8)	31.4 (+1.5)	31.2 (+1.3)	31.7 (+1.8)	31.7 (+1.8)
llama-3-8b	9.7	18.9 (-2.4)*	21.3	**23.8** (+2.5)	**19.3** (-2.0)	22.2 (+0.9)	**23.8** (+2.5)	**24.5** (+3.2)	22.6 (+1.3)	23.0 (+1.7)	23.4 (+2.1)
glm-4-9b	9.6	26.2 (-5.7)*	31.9	28.2 (-3.7)	**23.9** (-8.0)	**36.4** (+4.5)	**35.9** (+4.0)	**34.8** (+2.9)	**36.7** (+4.8)	**35.6** (+3.7)	**36.2** (+4.3)
llama-3-70b	10.3	**46.7** (+3.5)	43.2	44.6 (+1.4)	24.6 (-18.6)	42.8 (-0.4)	43.5 (+0.3)	43.6 (+0.4)	43.0 (-0.2)	43.2 (±0.0)	43.5 (+0.3)
qwen-2-72b	10.4	35.1 (-1.7)	36.8	37.4 (+0.6)	19.5 (-17.3)	37.2 (+0.4)	36.7 (-0.1)	36.7 (-0.1)	37.2 (+0.4)	37.3 (+0.5)	37.2 (+0.4)
qwen-1.5-110b	10.1	45.7 (+0.9)	44.8	42.8 (-2.0)	16.6 (-28.2)	44.7 (-0.4)	44.9 (+0.1)	44.6 (-0.2)	45.1 (+0.3)	45.0 (+0.2)	44.8 (±0.0)
gpt-3.5	9.9	**28.5** (+2.6)	25.9	27.1 (+1.2)	13.0 (-12.9)	26.5 (+0.6)	27.0 (+1.1)	**28.5** (+2.6)	26.5 (+0.6)	26.7 (+0.8)	27.2 (+1.3)
gpt-4	9.9	46.4 (-0.5)	46.9	46.9 (±0.0)	34.6 (-12.3)	47.3 (+0.4)	47.5 (+0.6)	47.7 (+0.8)	47.5 (+0.6)	47.8 (+0.9)	47.7 (+0.8)
ARC-Challenge											
mistral-7b	24.9	53.1 (-17.9)	71.0	70.3 (-0.7)	47.7 (-23.3)	71.7 (+0.7)	71.7 (+0.7)	71.6 (+0.6)	71.7 (+0.7)	71.7 (+0.7)	71.6 (+0.6)
llama-3-8b	25.2	44.4 (-21.8)	66.2	52.8 (-13.4)	41.1 (-25.1)	**71.3** (+5.1)	**70.0** (+3.8)	**70.0** (+3.8)	**71.6** (+5.4)	**71.3** (+5.1)	**71.3** (+5.1)
glm-4-9b	24.8	69.9 (-9.4)*	79.3	80.1 (+0.8)	65.1 (-14.2)	**82.7** (+3.4)	**82.3** (+3.0)	**82.0** (+2.7)	**82.8** (+3.5)	**83.0** (+3.7)	**82.7** (+3.4)
llama-3-70b	25.3	88.9 (+1.1)	87.8	87.9 (+0.1)	**80.8** (-7.0)	88.5 (+0.7)	88.4 (+0.6)	88.1 (+0.3)	88.5 (+0.7)	88.4 (+0.6)	88.4 (+0.6)
qwen-2-72b	24.8	84.7 (-1.1)	85.8	86.0 (+0.2)	36.7 (-49.1)	86.3 (+0.5)	86.2 (+1.3)	85.8 (±0.0)	86.3 (+0.5)	86.3 (+0.5)	86.2 (+0.4)
qwen-1.5-110b	24.7	87.0 (-0.7)	87.7	88.3 (+0.6)	53.4 (-34.3)	88.1 (+0.4)	88.1 (+0.4)	88.0 (+0.3)	88.1 (+0.4)	88.1 (+0.4)	88.1 (+0.4)
gpt-3.5	25.2	78.1 (+1.2)	76.9	77.0 (+0.1)	29.9 (-47.0)	78.2 (+1.3)	77.9 (+1.0)	78.2 (+1.3)	78.1 (+1.2)	77.9 (+1.0)	77.9 (+1.0)
gpt-4	25.0	92.9 (+0.4)	92.5	92.8 (+0.3)	**87.3** (-5.2)	92.9 (+0.4)	92.7 (+0.2)	92.8 (+0.3)	92.8 (+0.3)	92.8 (+0.3)	92.9 (+0.4)

Table 2: Overall accuracy results on MMLU, MMLU-Pro and ARC-Challenge benchmarks. 'Rand.' and 'Dicta.' denote 'random' and 'dictatorial', respectively. The numbers in parentheses are relative to the *blind dictatorial* baselines. Performance gains are marked in red, and loss in blue. Notable cases are marked in **bold**. *Results marked with asterisk are calculated utilizing partial profiles (see Appendix C).

Robustness against Unreliable Agents

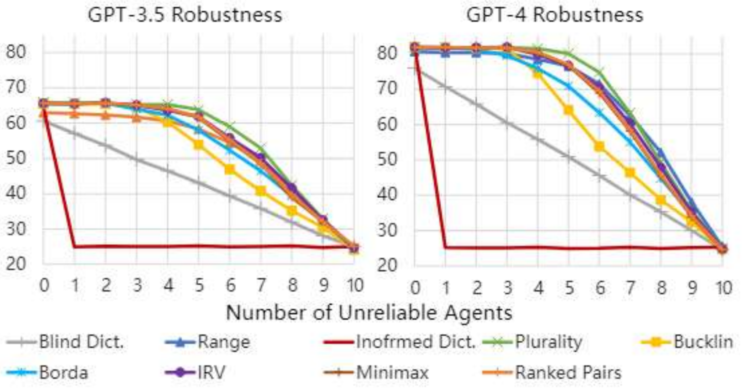

Figure 4: Accuracy impact of increasing number of unreliable agents built on gpt-3.5 and gpt-4.

Limitation：

- ➢ MCQA is a limited scenario of CDM (preference over correctness)
- ➢ Limited CDM methods in GEDI, no compound of multiple voting strategies
- ➢ Voting Tax: computation cost of inter-agent communication is high

VPPs in Australia

- **Thin Margins, High Competition**

Most energy service providers enter VPP market, i.e. installation of solar panels and battery storage systems, along with energy management software to monitor and control energy usage

- **Customer Experience**

Financial benefits, environmental impact and community benefits need to be clear

- **Operational Visibility, Dispatchability and Predictability**

Technical challenges in forecasting, orchestrating in VPPs in a highly complex cyber-physical-social system

- **Data Sharing Needs and Cybersecurity Threats**

VPPs open to new cybersecurity threat, when cloud-based solutions penetrate power system SCADA

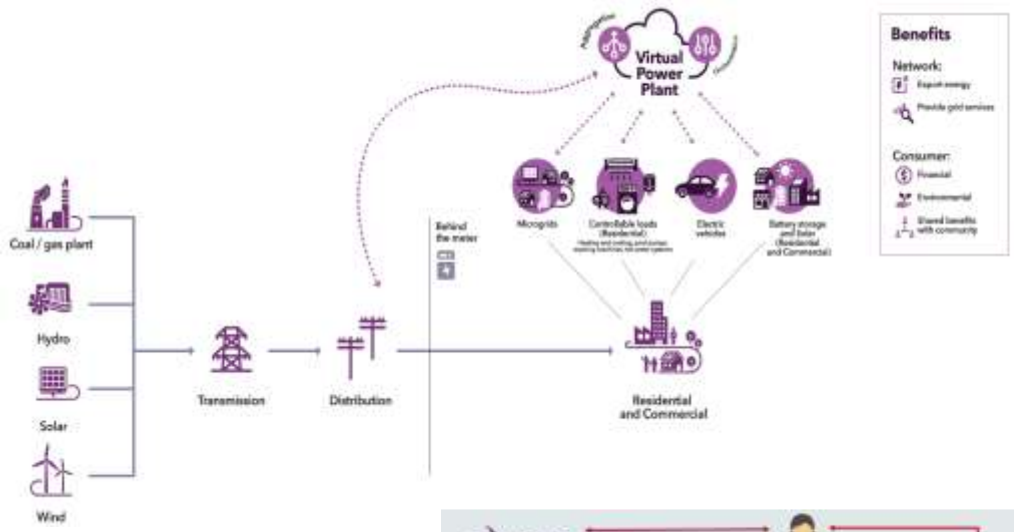

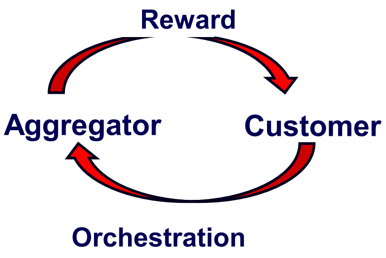

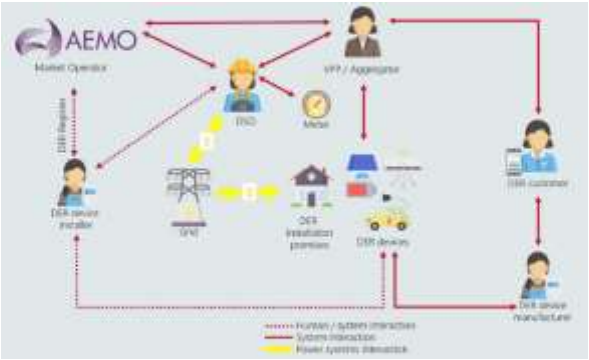

Source: AEMO NEM Virtual Power Plant Demonstrations Report, September 2021.

VPPs in Australia (cont.)

➢ **AEMO (Market System Operator)**

Market optimization and system security, visibility and coordination of DERs, modeling consumer VPP behaviors, forecast uncertainty from BTM DERs, compliance and registration, secure technical envelope, data exchange infrastructure

➢ **TNSP/DNSP (Network Service Provider)**

Enhance capability DERs, visibility monitor DER power flow, DOEs and DER planning, network and system optimization

➢ **Aggregators, Retailers**

Coordinating and optimize heterogeneous DERs to offer service under price signals and DOEs, complex dispatch and onboarding support, customer centric engagement

➢ **Consumers/Prosumers**

Low engagement due to unclear value proposition, confused on export limits, data privacy, difficult to set up and in monitoring

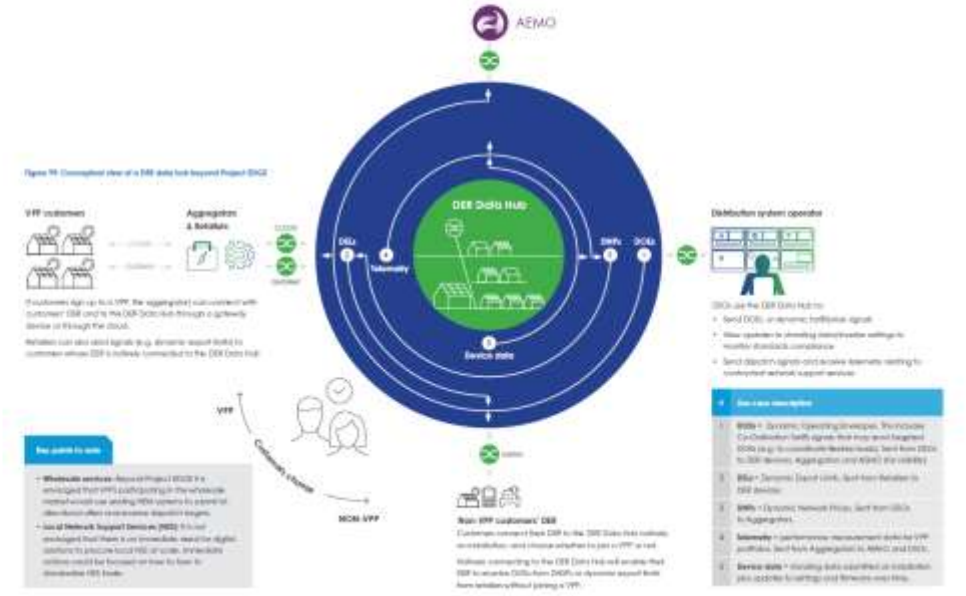

Project EDGE (Energy Demand and Generation Exchange): multi-year project to demonstrate an off-market, proof-of-concept Distributed Energy Resource (DER) Marketplace that efficiently operates DER to provide both wholesale and local network services within the constraints of the distribution network (AEMO, AusNet Services and Mondo, with financial support from the Australian Renewable Energy Agency (ARENA))

Potential GenAI Applications on VPPs

Beneficiary	GenAI Use Cases
Market Operator (AEMO)	• GPT to generate visualization and aggregated summaries from VPP data (DER visibility) • Agentic simulation on system-wide DER response, extreme dispatch schedules ,bidding behavior, DER flexibility (complexity modelling and coordination, secure envelop) • BTM DERs forecasting (LLM + time series + real world environmental data) • Automated compliance report generation (compliance) • AI-based protocol integration in data hub (DER interoperability)
DNSPs/TNSPs	• Agentic AI to simulate load profiles and DER adoption (DER planning, network optimization) • Visualize and simulate power flow (DERs monitoring and visualization) • Forecast congestion and generate dynamic limit recommendation (DER planning)
Aggregators and Retailers	• Agentic AI to generate adapters for APIs and heterogenous devices (DER interoperability) • Automate control scripts, automated bidding (DER coordination) • Forecast dynamic price and DER availability with historical and real time data (support DER coordination) • Agentic AI assistant for customer engagement, onboarding, help desk (customer engagement) • AI anomaly detection, root cause analysis/diagnostic, preventative maintenance (DER visibility) • GenAI to support design code check, design automation (DER efficiency)
Prosumers	• Agentic AI assistant guides or automates battery, EV charger, inverter configuration (customer experience) • Real-time explainable AI assistant to engage with customers on control rationality, export limits • Adaptive learning agents optimize scheduling of home appliance, EV/battery and solar panels based on price, demand, DOE, weather • Local GenAI running on edge devices to handle data privacy

RMIT Intelligent Informatics & Control Group (I^2C)

- Lead by Prof M Jalili, Dist. Prof X Yu with 6 staff members, 15+ research fellows and +40 HDRs
- Displine: Power Energy, Control, AI & Analytics, Industrial Application
- Industrial Partners: Jemena, Ausnet, AGL, Simens, Citipower/Powercor, Pacific Hydro, AUSTRC, Intyalheme, C4Net and more
- Research Funding: ARC ($6M+), Victoria Government ($6M+), CRC ($2M+), Industrial ($3M+)
- GenAI on DER coordination, charging scheduling, digital twin of distribution grid, EVs on grid, V2G, etc.

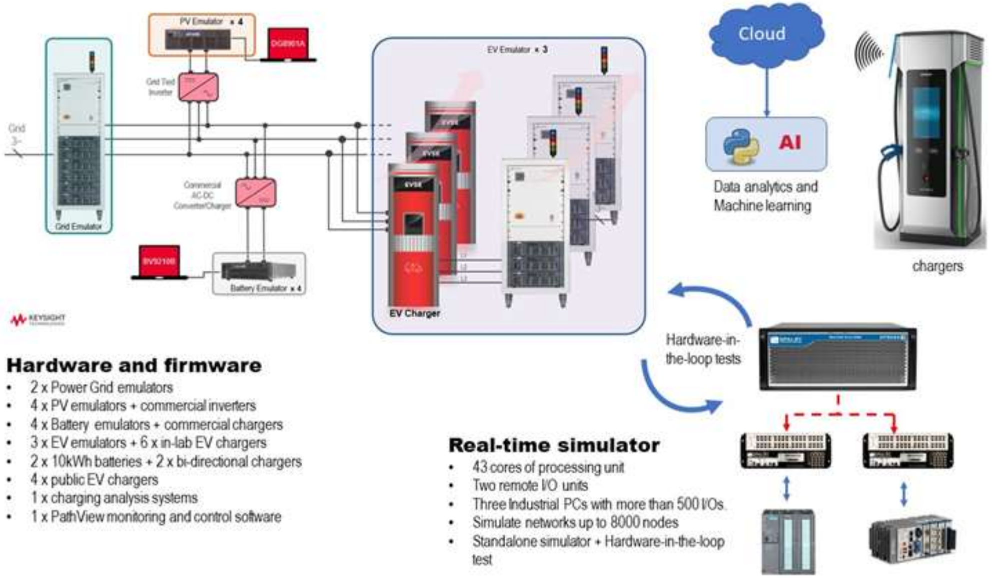

RMIT EV Living Lab

Major Literature

- Gartner, 2024, What's Driving the Hype Cycle for Generative AI, https://www.gartner.com/en/articles/hype-cycle-for-genai
- Gartner, 2024, Generative AI Use Cases Show Promise in Industry Segments, https://www.gartner.com/en/articles/generative-ai-use-cases
- McKinsey & Company Report, 2023, The Economic Potential of Generative AI: The next productivity frontier
- John, Searle. 1980. Minds, Brains and Programs. *Behavioral and Brain Sciences*, 3: 417-457.
- Stevan, Harnad. 1990. The Symbol Grounding Problem. *Physica D: Nonlinear Phenomena*, 42(1-3): 335–346.
- Rodney A. Brooks. 1991. Intelligence without Representation. *Artificial Intelligence,* 47:139- 159.
- David J. Chalmers. 1992. Subsymbolic Computation and the Chinese Room. In *Dinsmore, J. (ed.) The Symbolic and Connectionist Paradigms: Closing the Gap*. Hillsdale, NJ: Lawrence Erlbaum.
- Georg Dorffner and Erich Prem. 1993. Connectionism, Symbol Grounding, and Autonomous Agents. In P*roceedings of the Fifteenth Annual Meeting of the Cognitive Science Society*, pages 144-148. Hillsdale, NJ: Lawrence Erlbaum.
- Richard Sutton, 2024, New Path for AI: Approximately Correct Podcast, URL: https://www.youtube.com/watch?v=NvfK1TkXmOQ
- Ranjan Sapkota, Konstantinos I. Roumeliotis, Manoj Karkee, 2025, AI Agents vs. Agentic AI: A Conceptual Taxonomy, Applications and Challenges, https://arxiv.org/abs/2505.10468
- Xiao Liu, Hao Yu, Hanchen Zhang, Yifan Xu, Xuanyu Lei, Hanyu Lai, Yu Gu, Hangliang Ding, Kaiwen Men, Kejuan Yang, et al. Agentbench: Evaluating llms as agents. ArXiv preprint, abs/2308.03688, 2023. URL https://arxiv.org/abs/2308.03688.
- Aohan Zeng, Mingdao Liu, Rui Lu, Bowen Wang, Xiao Liu, Yuxiao Dong, and Jie Tang. 2024. AgentTuning: Enabling Generalized Agent Abilities for LLMs. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 3053–3077, Bangkok, Thailand. Association for Computational Linguistics.
- Yifan Song, Da Yin, Xiang Yue, Jie Huang, Sujian Li, and Bill Yuchen Lin. 2024. Trial and Error: Exploration-Based Trajectory Optimization of LLM Agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7584–7600, Bangkok, Thailand. Association for Computational Linguistics.
- Yifan Song, Weimin Xiong, Xiutian Zhao, Dawei Zhu, Wenhao Wu, Ke Wang, Cheng Li, Wei Peng, and Sujian Li. 2024. AgentBank: Towards Generalized LLM Agents via Fine-Tuning on 50000+ Interaction Trajectories. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 2124–2141, Miami, Florida, USA. Association for Computational Linguistics.
- Weimin Xiong, Yifan Song, Xiutian Zhao, Wenhao Wu, Xun Wang, Ke Wang, Cheng Li, Wei Peng, and Sujian Li. 2024. Watch Every Step! LLM Agent Learning via Iterative Step-level Process Refinement. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1556–1572, Miami, Florida, USA. Association for Computational Linguistics.
- Xiutian Zhao, Ke Wang, and Wei Peng. 2024. An Electoral Approach to Diversify LLM-based Multi-Agent Collective Decision-Making. *In Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 2712–2727, Miami, Florida, USA. Association for Computational Linguistics.
- AEMO, 2021, AEMO NEM Virtual Power Plant Demonstrations Report, https://arena.gov.au/assets/2021/09/aemo-virtual-power-plant-demonstrations-report-4.pdf
- Project EDGE, 2023, Project EDGE Final Report, https://aemo.com.au/-/media/files/initiatives/der/2023/project-edge-final-report.pdf?la=en

Thanks
Questions & discussion welcome.