Agentic AI Systems Reading Group

Mentors: Lunyiu Nie (flitternie@gmail.com)

Description

Our group explores agentic AI systems—agents that act, communicate with humans and other agents, and ground their behavior in real-world contexts. We progress from core LLM foundations to tool use, planning/acting, multimodal grounding, multi-agent orchestration, evaluation, and safety.

Students will read and present every week (2–3 papers), with discussion focused on mechanisms, grounding, and robust deployment. We will also host invited talks by PhD students working in relevant research areas; details will be announced.

Prerequisites

Audience: CS undergraduates with Python Programming experience.
Recommended: Introductory NLP.

Resources

- https://agenticai-learning.org/sp25
- https://agenticai-learning.org/f25
- https://esteng.github.io/grounded-communication/

Proposed Agenda (Subject to change)

Week	Content Covered
Week of September 15th	**Foundations: Transformer era + emergent reasoning** *Core mechanics of LLMs and evidence for reasoning behaviors.* - Attention Is All You Need (Vaswani et al., 2017) - Language Models are Few-Shot Learners (GPT-3) (Brown et al., 2020)

	- Chain-of-Thought Prompting Elicits Reasoning (Wei et al., 2022)
Week of September 22nd	**From language to action: semantic parsing & tool use** *From text → executable structure and tools.* - Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task - Toolformer: Language Models Can Teach Themselves to Use Tools (Schick et al., 2023) - - ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs
Week of September 29th	**Programs as skills: induction & abstraction** *Moves from single calls to structured, reusable skills*. - DreamCoder: Growing generalizable, interpretable knowledge (Ellis et al., 2021) - Voyager: An Open-Ended Embodied Agent with LLMs (Wang et al., 2023) - ReGAL / Refactoring Programs to Discover Generalizable Abstractions (Chen et al., 2023)
Week of October 6th	**Computer-use agents: browsers, apps & evaluation at scale** High-variance, open-world tasks and robust evals. - Mind2Web: Towards a Generalist Agent for the Web (Deng et al., 2023) - WebArena: A Realistic Web Environment for Autonomous Agents (Zhou et al., 2023) - - An Illusion of Progress? Assessing the Current State of Web Agents (2024)
Week of October 13th	**Orchestrating many agents: specialization & division of labor** *Pipelines and modularization.* - AutoGen: Multi-Agent Conversation for LLM Apps (Wu et al., 2023) - Predicting Multi-Agent Specialization via Task Parallelizability (2024) - Resource-efficient Inference with Foundation Model Programs (2025)
Week of October 20th	**Multi-agent consensus: debate & mixture-of-experts** *Adds social reasoning and ensemble-style control.*

	- Improving Factuality & Reasoning via Multi-Agent Debate (Du et al., 2023) - Encouraging Divergent Thinking via Multi-Agent Debate (Chan et al., 2023) - ReConcile: Round-Table Consensus among Diverse LLMs (2024)
Week of October 27th	**Vision–language foundations for grounding** *Ground percepts before embodied control.* - CLIP: Learning Transferable Visual Models from Natural Language (Radford et al., 2021) - BLIP-2 (Li et al., 2023) - ViperGPT & VisProg
Week of November 3rd	**Learning to act with tools: optimization & RL for agents** *Introduces optimization signals for tool-using agents.* - Gorilla: LLM Connected with Massive APIs (Patil et al., 2023) - Agentic Reasoning and Tool Integration via Reinforcement Learning (2024) - ToolRL: Reward is All Tool Learning Needs (2024)
Week of November 10th	**LLM Agents for Math** - AlphaProof: when reinforcement learning meets formal mathematics - Draft, Sketch, and Prove: Guiding Formal Theorem Provers with Informal Proofs - An In-Context Learning Agent for Formal Theorem-Proving
Week of November 17th	**Agentic Systems Safety** - DataSentinel: A Game-Theoretic Detection of Prompt Injection Attacks - AgentPoison: Red-teaming LLM Agents via Poisoning Memory or Knowledge Bases - Progent: Programmable Privilege Control for LLM Agents
Week of December 1st	TBC