

Large Language Models

(Part 4)

April 21, 2026

Announcements

- + Final Project Presentations on April 23 and 28
 - + 8-9 minutes long (typically, 1-1.5 minutes per content slide)
 - + Remember that we are not experts in your subject matter, so limit the use of subject-specific jargon
- + Final Code + Report due May 7 at 11:59pm ET
- + I will need to convert dsip-s26 into a private repository after the semester, so be on the lookout for an invitation to (re-)join this repository.

Plan for today

+ **Agentic AI**

Agentic AI

What is Agentic AI?

- + **Agentic AI is an AI that “acts”**
- + While traditional generative AI (like a standard chatbot) waits for a prompt to output text/images, agentic AI is designed to be autonomous and can create multi-step plans and use external tools to get some task done
- + What makes an AI “agentic” is a loop of four specific behaviors:
 - + **Reasoning & Planning:** It doesn't just guess the next word; it breaks a complex goal (e.g., "Organize a team lunch") into logical steps.
 - + **Tool Use:** It can “play well” with other software. It can browse the web, execute code, or call APIs to perform real-world actions.
 - + **Autonomy:** It operates with minimal human supervision. Once you give it a high-level objective, it manages the "how" on its own.
 - + **Iterative Correction:** If a step fails (e.g., a restaurant is fully booked), it doesn't just stop. It realizes the error, "replans," and tries an alternative.
- + This field is rapidly evolving!

What is Agentic AI?

Real-World Examples

- + **Customer Support:** Instead of just answering "What is your refund policy?", an agentic system checks the user's order history, verifies the return window, generates a shipping label, and updates the warehouse system.
- + **Software Development:** "Coding agents" can now find a bug, write the fix, run tests to make sure it works, and submit a pull request for a human to review.
- + **Finance:** Systems at firms like JPMorgan Chase now use agents to monitor millions of transactions in real-time, autonomously flagging and blocking fraud patterns as they evolve, rather than waiting for a human to update the rules.
- + **Personal Productivity:** A "travel agent" AI that doesn't just suggest flights but actually navigates the booking sites, handles the payment (with permission), and adds the digital ticket to your mobile wallet.

Agentic AI in Action

Navigate to dsip-s26/ course repository

- + `git pull`
- + Install and activate conda environment
- + Open `course_materials/llms/llm_agents.ipynb`

Agentic AI in Action
