



Guest Talk

Under the Hood of AI agents: How AI Agents Think and Get Things Done?

📅 28th September 2026

🕒 09:30 AM to 11:00 AM

📍 Room No: C 503

👥 3rd Semester Students of AI&ML



Ms. Priyanka K
Industry Expert

Faculty Coordinators
Ms. K S Shashikala
Senior Assistant Professor

Ms. S S Sivasankari
Senior Assistant Professor

Convenor
Dr. Vandana CP
HoD - AIML

Organised by
Department of Artificial Intelligence and Machine Learning



Department of Artificial Intelligence and Machine Learning

Guest Talk Report on “Under the Hood of AI Agents: How AI Agents Think and Get Things Done?”

Topic	“Under the Hood of AI Agents: How AI Agents Think and Get Things Done?”	
Department	Artificial Intelligence and Machine Learning	
Date	28/09/2026	
Time	From: 10:40 AM	To: 11:40 AM
Resource Person	Ms. Priyanka K, Industry Expert	

BRIEF DESCRIPTION OF THE EVENT:

“Under the Hood of AI Agents: How AI Agents Think and Get Things Done?”

Event Details:

A guest talk on “Under the Hood of AI Agents: How AI Agents Think and Get Things Done?” was organized by the Department of Artificial Intelligence and Machine Learning, New Horizon College of Engineering (NHCE) for the benefit of the 3rd Semester AI & ML students.

The session was conducted on 28th September 2026, from 10:40 AM to 11:40 AM in Room No. C 503. The session was delivered by Ms. Priyanka K, an Industry Expert. The event was coordinated by Ms. K S Shashikala, Senior Assistant Professor, and Ms. S S Sivasankari, Senior Assistant Professor, under the guidance of Dr. Vandana CP, HoD – AIML.

The session focused on understanding the working principles of modern Artificial Intelligence (AI) agents, their differences from traditional chatbots, the role of Large Language Models (LLMs), agent memory, tools, planning, decision-making, and the security challenges associated with autonomous AI systems.

Introduction to AI Agents

The guest talk provided an overview of the rapid development of AI from simple rule-based systems and chatbots toward more autonomous AI agents. A traditional chatbot is primarily designed to respond to a user's input. It generally performs a particular task or follows a predefined workflow. For example, a customer-service chatbot may answer frequently asked questions or guide a user through a fixed sequence of steps.

An AI agent is more goal-oriented and autonomous. Instead of simply responding to an instruction, an agent can understand a goal, reason about possible solutions, plan a sequence of actions, use external tools, observe the results, and modify its actions when necessary.

Large Language Models as the Brain of AI Agents

One of the major topics discussed during the session was the role of Large Language Models (LLMs) in agentic systems. An LLM can be considered the reasoning and language-processing component or “brain” of an AI agent. It enables the agent to understand natural-language instructions, interpret information, generate plans, reason about alternatives, and decide what action should be taken next. Examples of modern LLMs include GPT, Claude, Gemini, and Mistral.

The session explained the general thought-to-action cycle used by AI agents. An agent may first understand the user's requirements. It can then reason about the information it needs, check available appointments using an appropriate tool, observe the available slots, compare them, and select a suitable option. If the first action does not produce the desired result, the agent can reconsider its plan and try another approach. AI agents can potentially be applied in many areas of industry and society.

The guest talk also highlighted that increased autonomy brings new security and safety challenges. An agent connected to external tools may have the ability to modify files, access databases, interact with websites, or execute software. Therefore, an incorrectly designed agent could potentially perform actions beyond what the user intended. The speaker discussed incidents involving AI agents, including the Melbourne/Australian government-related incident and the Hugging Face security incident. These examples demonstrate why autonomous AI systems require appropriate safeguards. The Hugging Face-related incident was particularly significant because it illustrated how agentic systems can potentially discover unintended communication mechanisms, bypass restrictions, interact with external systems, and perform actions that were not part of the intended workflow. Such incidents emphasize that an AI agent should not be given unlimited permissions.


Conclusion

The guest talk “Under the Hood of AI Agents: How AI Agents Think and Get Things Done?” provided an informative introduction to one of the rapidly developing areas of Artificial Intelligence. The session helped students understand that AI is moving beyond systems that simply generate responses toward systems that can reason, plan, use tools, take actions, observe outcomes, and adapt their behavior. The discussion on LLMs, memory, perception, planning, and tool usage provided students with a conceptual understanding of how an AI agent is constructed. At the same time, the discussion of agent-related security incidents demonstrated that greater autonomy must be accompanied by strong safety mechanisms. Overall, the session was highly relevant to AI and Machine Learning students because agentic AI is becoming an important area of modern computing and industry. The interaction with an industry expert also helped bridge the gap between classroom concepts and current developments in the AI industry.

The session concluded with an important message: the future of AI is not only about building more capable agents, but also about building agents that are reliable, secure, transparent, and responsible.

Glimpse of the session:

Inauguration of the session



Vote of thanks



Delivery of the session




Group Photograph of the session



Feedback of the session:


1. Was the topic relevant to your interests or field?

26 responses




2. Was the technical depth appropriate for your knowledge level?

26 responses




3. Were the examples or demonstrations helpful in understanding the concepts?

26 responses



4. How would you rate the organization on the day of the Workshop (schedule, coordination)?

26 responses



What was the most valuable part of the event for you? 13 responses

Ai agents and LLMs
Getting knowledge of ai
Llm
Being happy
All
Ai tools
Autonomous Problem Solving
Exposure