Conversational Multi-Agent System For Smart Campus Energy Management

IMPERIAL
Ada Lovelace Academy
FACULTY OF ENGINEERING

By
Arastun Mammadli
Supervisors:
Prof. Lee Stott (Microsoft)
Prof. Gerard Gorman (Imperial College London)

Abstract

Large Language Models (LLMs) are increasingly used in autonomous multi-agent systems. They have seen extensive success in the fields of web navigation, software engineering, personalised learning and research. However, due to its technical and multi-facet requirements (dynamic control and monitoring, feedback collection, sustainability awareness), energy management remains a challenging domain. An agentic system holds potential to address these challenges, serving a range of campus stakeholders (students, faculty, and administrators). This project leverages emerging frameworks in Microsoft's AI Agent Services to demonstrate a proof of concept for a conversational campus energy management system. Building on base agents, it integrates language and search services to enable a robust orchestration. The system includes a user-facing chatbot that can respond to system-related queries, collect student feedback, and assist administrators with energy monitoring prognostics and chart-based visualisations. The project also designs a synthetic energy data schema with hierarchical infrastructure, usage categories, and environmental context that mirrors real world systems and supports agentic capabilities. The system is evaluated on modular (individual agents and services) and end-to-end levels. Special focus is given to explore intent routing, access-level based grounding mechanisms, and query complexity. All together the assessments outline practical challenges of composite agentic workflows and their integration with external services. They discuss promising future directions through integrations with real-world energy systems.

32%
of global energy demand is from the building sector
(UNEP/GlobalABC, 2025)

Students and Staff
Provide feedback, learn about initiatives

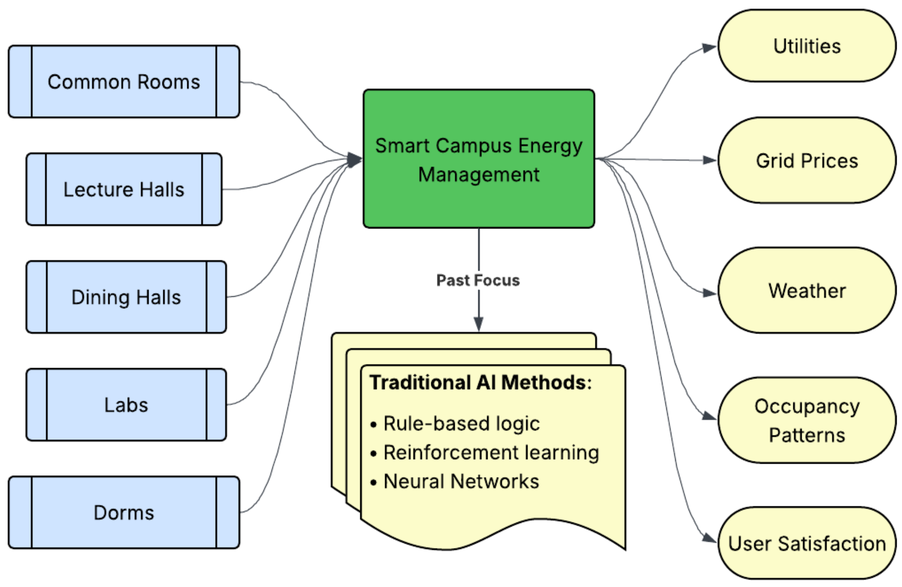

Administrators
Analysis, prognostics, textual and visual summaries

The Problem

University campuses serve a range of stakeholders and have distinct energy profiles and usage patterns. Traditional systems are rigid and custom-made.

Motivation

Universities have a strong drive towards sustainability, with many institutions committing to net-zero emission targets. Imperial College London aims to be net-zero by 2040 (Imperial College London, 2024).

Potential

Past works have mostly explored multi-agent deep reinforcement learning (MADRL) techniques (Yu et al., 2021). They are energy monitoring only solutions, with limited user interaction and adaptability.

There is an opportunity to explore emerging concepts and frameworks in the context of campus energy management. To address the complex institutional needs, while serving a range of stakeholders (students, faculty, and administrators).

System Design

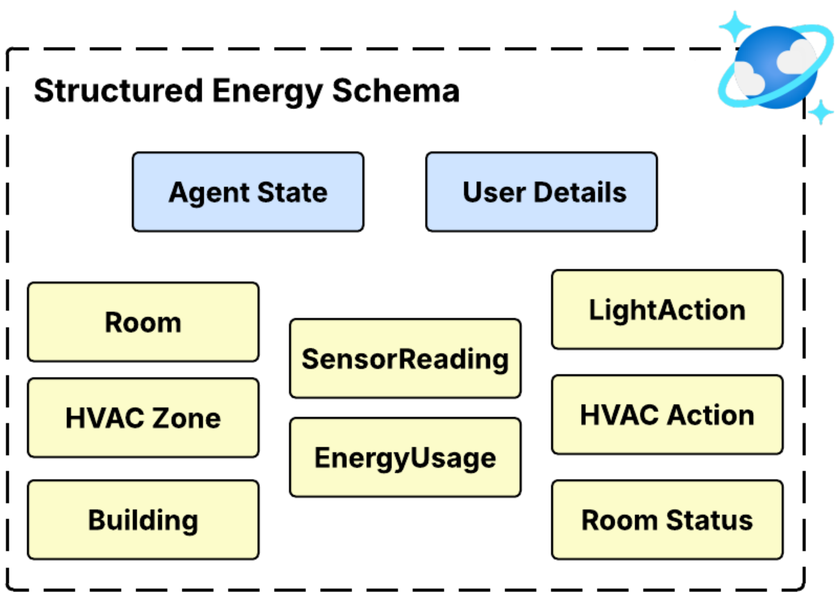

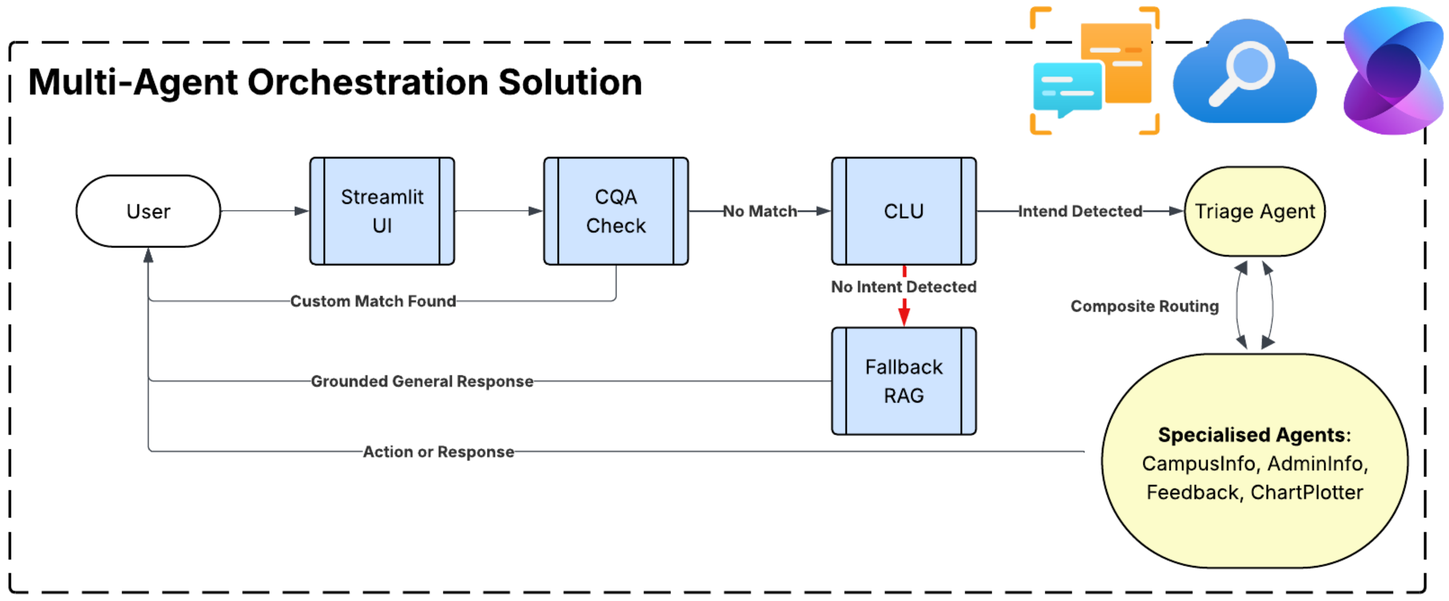

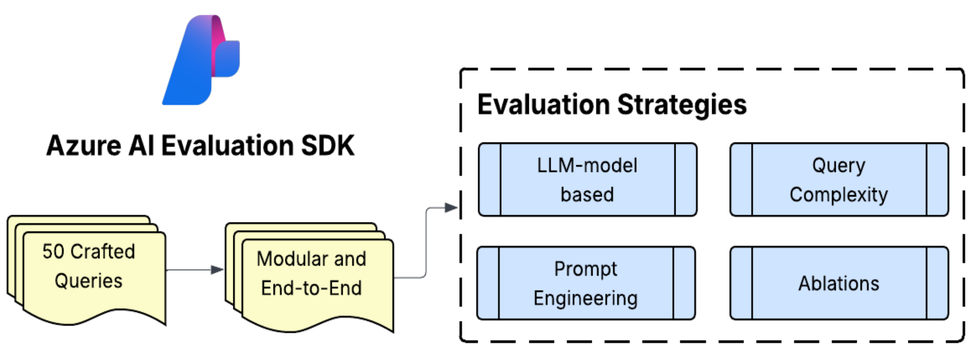

Key Deliverables

- Composite agentic workflow with 4 specialised agents (CampusInfo, AdminInfo, Feedback, ChartPlotter) integrated with Language and Search services, orchestrated through a Triage agent.
- Structured energy data schema built on real-world datasets and monitoring systems (Moraliyage et al., 2022). Heterogeneous set of containers that can be used to comprehensively test the agentic capabilities.
- A comprehensive evaluation pipeline to outline the challenges of composite agentic workflows and their integration with external services.

ReAct Style Few-Shot Exemplar

Question: Can you give me chart of occupancy history in the room Lecture Theatre 1?
Thought 1: I first need to determine if Lecture Theatre 1 exists, and get its id, I can then use plot_room_status as the user is asking about occupancy status, and return the chart and chart summary to the user.
Act 1: Call[get_room_id(Lecture Theatre 1)]
Observation 1: the room Lecture Theatre 1 exists with id LT1
Thought 2: Now I can call plot_room_status to generate the chart and obtain its file path.
Act 2: Call[plot_room_status(LT1, occupancy)]
Observation 2: I get the generated figure for occupancy history
Act 3: Finish[{"text": <Here is the occupancy history chart from Lecture Theatre 1>, "chart": [<"path_to_plot">]}]

Performance Analysis & User Experience

Composite Intent Routing

To improve intent resolution the project employs language services. It incorporates Conversational Language Understanding (CLU) and grounds the Triage with classified intent and confidence score.

The Triage agent acts as a coordinator and manages composite requests to specialised agents (CampusInfo, AdminInfo, Feedback, and ChartPlotter) in a multi-agent group chat setting.

Ablation studies show that incorporating a Triage agent significantly improves response quality. It leads to more concise responses while maintaining the performance score.

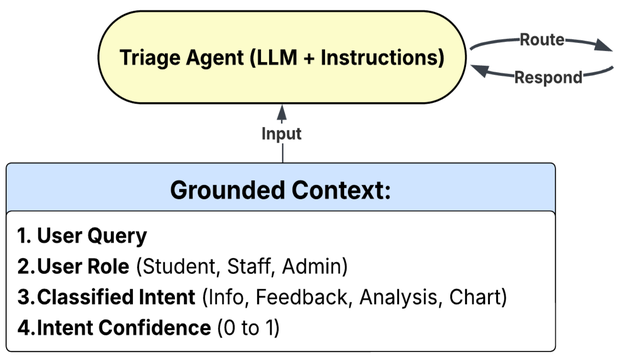

Multi-Step Reasoning

We explore multi-step reasoning capabilities by evaluating with single and multi-utterance samples.

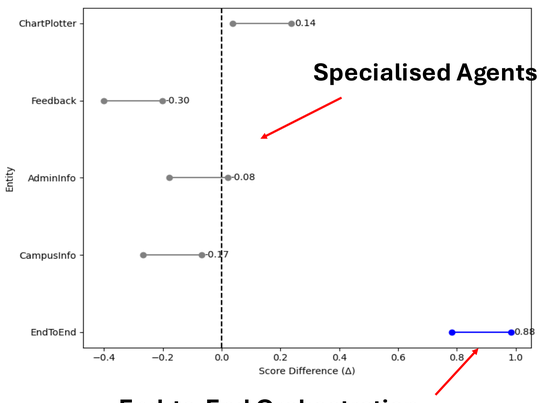

Agents or the Orchestration?

At end-to-end level, the system handles single-utterance queries much better than multi-utterance ones (Δ = 0.88). This gap is not present at modular agent-only evaluations.
Suggesting coordination and routing remain key bottlenecks.

Prompt Engineering

We find that ReAct-style (Yao et al., 2023) prompting affects response length: longer outputs for information-retrieval-based agents, shorter for action agents—reflecting step-by-step reasoning vs. concise action planning.

ChartPlotter outputs visual images (with captions) and therefore unaffected.

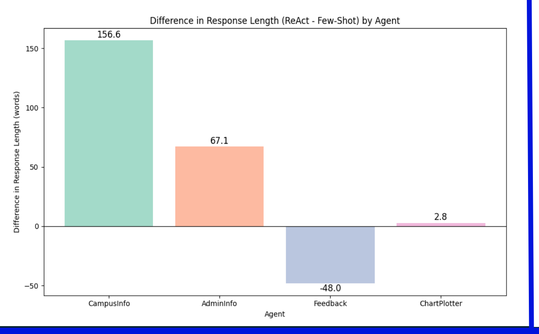

User Experience

Admins have higher-access level to information stored by the campus energy dataset (e.g., registered user details, sensor readings, live room status, past energy usage logs). They can request system information, visual prognostics, user feedback summary, etc.

Students and faculty have restricted access. They can request personal details, general information (sustainability goals, energy initiatives). They can also provide live feedback for their room.

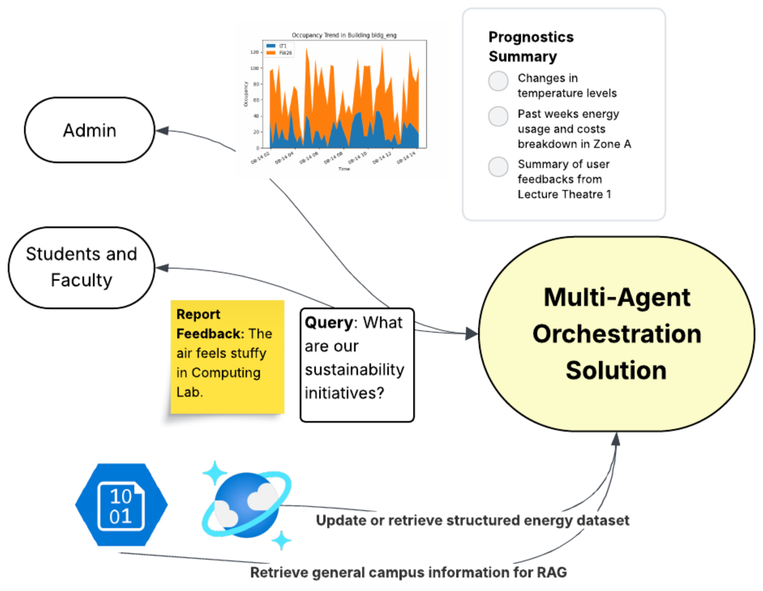

References
UN Environment Programme and Global Alliance for Buildings and Construction (2025) Global status report for buildings and construction 2024/2025. Imperial College London (2024). *Annual Sustainability Report 2023-24*. Yu, L., Qin, S., Zhang, M., Shen, C., Jiang, T., & Guan, X. (2021). A review of deep reinforcement learning for smart building energy management. *IEEE Internet of Things Journal*, 8(15), 12046–12063. Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2023). React: Synergizing reasoning and acting in language models. *International Conference on Learning Representations*. Moraliyage, H., Mills, N., Rathnayake, P., De Silva, D., & Jennings, A. (2022). Unicon: An open dataset of electricity, gas and water consumption in a large multi-campus university setting.