AmI HMAS: A Hypermedia MAS for Goal-Driven Interactions with every-day Smart Environments

Demonstration Track

Alexandru Sorici
POLITEHNICA Bucharest
Bucharest, Romania
alexandru.sorici@upb.ro

Victor-Vasile Udriștioiu
POLITEHNICA Bucharest
Bucharest, Romania
victor.udristioiu@stud.acs.upb.ro

Ioan-Cristian Cordoș
POLITEHNICA Bucharest
Bucharest, Romania
ioan.cordos@stud.acs.upb.ro

Andrei Olaru
POLITEHNICA Bucharest
Bucharest, Romania
andrei.olaru@upb.ro

ABSTRACT

Ad-hoc, goal-driven interactions in smart environments (homes, offices, hotels) have been a long lasting objective in Ambient Intelligence (AmI). Advances in Large Language Models (LLM) for reasoning and use of hypermedia environments for multi-agent systems are bringing this objective closer to achievement. We demonstrate AmI HMAS, a framework for goal-driven, LLM-supported interactions with smart environments. AmI HMAS maps existing HomeAssistant deployments into semantically represented, navigable Hypermedia Environments, enabling discovery of real-world smart devices. The framework combines classic agency with LLM reasoning to perform environment exploration, request interpretation, and action planning, while also leveraging an engine that enables storage and reuse of past interaction experiences during reasoning. We showcase an end-to-end workflow, including environment mapping, user request handling, LLM reasoning, and the reuse of learned plans for effective smart environment interactions.

KEYWORDS

Hypermedia MAS, Web-of-Things, Ambient Intelligence, AI Agents, Goal-driven Interaction

ACM Reference Format:
Alexandru Sorici, Victor-Vasile Udriștioiu, Ioan-Cristian Cordoș, and Andrei Olaru. 2026. AmI HMAS: A Hypermedia MAS for Goal-Driven Interactions with every-day Smart Environments: Demonstration Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https: //doi.org/10.65109/MAXW5671

1 INTRODUCTION AND APPLICATION DOMAIN

Goal-driven and *assisted* interactions with smart environments (e.g. residential home, office, hotel) have been a long term objective in Ambient Intelligence (AmI) . Fluid interactions and integrated or discoverable services are still far from those envisioned in the 2001 ISTAG report (e.g. the Maria scenario [11]). Advances in web engineering and services bridging the physical and the digital world (e.g. IFTTT [3], HomeAssistant [5], openHAB [6], Google Home) are bringing the objective closer to practical use.

However, such services still have challenges to overcome. For example, Google Home interactions require *explicit* indication of a *single* action on a *single* device; HomeAssistant recipes provide device choreography, but are hard to define for non tech-savvy users and must be programmed in advance; IFTTT is more user friendly, but mostly involves simple rules that cover a triggering condition and an execution on a single device / service.

Yet, desired interactions require the ability to define *ad-hoc*, *immediate effect* or *persistent*, *implicitly*- (e.g. "The air feels stuffy in here", commanding a bath in the ISTAG Maria scenario) or *explicitly*-phrased interactions (e.g. "Bring the air humidifier on at 50% for 10 minutes after the oven has finished cooking"). Acknowledging this need, several recent works (e.g. Sasha [12], HomeBench [13], SimuHome [16]) introduce benchmarks of smart environment interactions focusing on several challenges frequently coming up in such ad-hoc, daily-use scenarios: correctly identifying which *affordances* of smart devices are called on through natural language, filtering out impossible (e.g. due to lack of capabilities) or ineffective (e.g. incorrect value ranges, asking for state change when the state is already correct) requests, dealing with immediate or maintenance goals, handling implicit or vaguely specified requests.

The benchmarks highlight another relevant aspect of recent research work: the increased potential for *understanding* and *reasoning about* natural language requests enabled by advances in AI foundational models, notably Large Language Models (LLMs). Such models are tested as means for interfacing with users (understanding requests and presenting resulting solutions), as well as the planning itself [10, 17].

However, one significant limitation of the mentioned smart environment simulators and the types of evaluation they facilitate is that the benchmark interactions are considered in isolation from each other. Moreover, the support environments are created in a simple, ad-hoc manner without considering the problem of facilitating *discovery*, *representation* and *management* of existing, real-world deployed smart environments.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/MAXW5671

The framework we are demonstrating[1], called AmI HMAS[2], has the following contributions: **(i)** an engine to map existing HomeAssistant [5] deployments to semantically represented (RDF-based), navigable hypermedia environments [8]; **(ii)** use of a signification-based mechanism and methodology to *remember/reuse* and *learn from* past interactions; **(iii)** development of LLM-supported agent *roles* that classify and interpret user requests, monitor environment capabilities and state changes, plan solutions for immediate or persistent user goals leveraging past interactions, and keep the user in the loop for plan validation.

2 TECHNOLOGY AND DEMONSTRATION

AmI HMAS is a framework for goal-driven interactions with smart environments that combines three agent and environment development paradigms into its functionality: ThingDescription (TD [4]) hypermedia-based agent environments (HMAS [8, 9], agent role-based organization of functionality using reactive and cognitive behaviors (by means of the SPADE agent development framework [7]), and LLM-supported reasoning for request understanding and plan generation.

2.1 AmI HMAS Framework

AmI HMAS agents are designed to operate in a TD-based hypermedia environment that follows the design principles of the Agents & Artifacts paradigm [15]. **The HomeAssistant integration engine** maps a deployed HomeAssistant instance to an HMAS environment. First, the adapter maps HomeAssistant *area_ids* to *workspaces* (conceptual environment functionality containers), treating each physical area as a separate workspace containing multiple *artifacts* (devices, sensors, services). HomeAssistant devices are then exposed as W3C TD Artifacts in RDF/Turtle format, with each device becoming a uniquely addressable artifact within its workspace. HomeAssistant Entity attributes and state values are exposed as readable *Property Affordances*, with special handling for *state* (current value) and *metadata* (timestamp of last change and other diagnostics). Entity services are mapped to *Action Affordances* with endpoints that define an explicit JSON schema for input parameters and action invocation output. State changes are pushed to any external monitors (including the AmI agent roles) using WebSub event notifications, forwarding HomeAssistant's real-time updates as RDF-based payloads.

The AmI HMAS agents are configured to *discover* a mapped HMAS environment following the W3C Web-of-Things Discovery protocol [2]. There are three main agent roles that provide the functionality of AmI HMAS. The ***EnvExplorer*** is the agent that *discovers* and catalogs the environment, indexing the *affordances* of available artifacts and tracking changes in PropertyAffordance values. The agent also manages an engine that records and recalls past affordance *usage experiences*. The ***UserAssistant*** receives user requests and determines their feasibility, based on capabilities known from the *EnvExplorer*. Feasible requests are classified by type of request and passed for solving to the *InteractionSolver*. Returned action invocation plans are presented in a user-understandable manner for confirmation. Confirmed plans are marked for storage as experience and their execution life-cycle is managed by the *UserAssistant*. The ***InteractionSolver*** agent is the main reasoning component, creating the plan of *ActionAffordance* invocations solving the request. *InteractionSolver* planning is supported by LLM-based reasoning. To support this, the *InteractionSolver* will ask *EnvExplorer* to provide *signifiers* for past experiences with identical or similar requests, so that planning considers only affordances that are relevant for the current request. **The Signifier Memory Engine** operates as a unit that facilitates the *reuse* of past interaction experiences. It relies on the concept of *signification*, borrowed from Affordance Theory [14]. In the AmI HMAS framework, a *signifier* binds an artifact *affordance* to an *intent* and a *context of execution*. Signifiers are represented in a structured manner, using RDF and SHACL [1] (for context conditions). The matching of a new request to existing signifiers takes into account *intent* similarity, as well as *context* similarity. The functionality of the AmI HMAS agents and the signifier representation has been described in more detail in previous work [18] and is further exemplified in the demo video.

2.2 Demonstrator

The demonstrator presents an end-to-end interaction workflow that integrates a HomeAssistant deployment with virtual devices, its hypermedia-based semantic representation, and the AmI HMAS framework agents supporting both explicit and implicit user requests. The HomeAssistant deployment includes a smart light, motorized blinds, and indoor / outdoor environmental sensors.

With HomeAssistant running, the mapping engine automatically translates the configuration into a TD-based Hypermedia Environment. The resulting RDF model is inspected through a web browser, illustrating how devices, states, and affordances are exposed as navigable semantic resources. When AmI HMAS agents start, logs illustrate agent exploration of the hypermedia environment and discovery of available artifacts. The demo then shows *explicit* user-driven interactions while highlighting the collaboration of agent roles: checking for the smart light state, LLM planning for multi-command requests to adjust light and blinds, storing the experience as signifiers mapping a goal to concrete affordances.

Experience reuse is demoed for *implicit* requests. With the lab in low light conditions, two solution pathways – an LLM-based environment exploration and planning for implicit requests and a *fast-path* reuse of stored Signifiers as *affordance hints* and plan adaptation – are presented when the user complains about insufficient illumination in two differently phrased requests. In supplementary material [3] we additionally discuss qualitative and quantitative results of this scenario, looking at planning method, success rates and latency, as well as signifier reuse rates.

3 CONCLUSION

AmI HMAS is a relevant example of successfully fitting together classical agency (supported through mature frameworks such as SPADE), frontier advances in agentic AI workflows and hypermedia-based agent environment modeling, applied to the domain of smart environments, which takes full advantage of the best of these developments and has the potential of real end-user impact.

[1]Demo video link: https://youtu.be/qPZIZ1Rz6eY
[2]Demo code: https://github.com/aimas-upb/llm-agents-for-ami/tree/aamas2026demo
[3]https://tinyurl.com/AmI-HMAS-AAMAS2026-sup

ACKNOWLEDGMENTS

This research was supported by the project "Romanian Hub for Artificial Intelligence - HRIA", Smart Growth, Digitization and Financial Instruments Program, 2021-2027, MySMIS no. 351416.

REFERENCES

[1] 2017. SHACL W3C Recommendation. https://www.w3.org/TR/shacl/. Accessed: 2026-01-05.

[2] 2023. W3C Web-of-Things Discovery Protocol Specification. https://www.w3.org/TR/wot-discovery/. Accessed: 2026-01-05.

[3] 2025. If-This-Than-That home automation service. https://ifttt.com/. Accessed: 2026-01-05.

[4] 2025. Web of Things (WoT) Thing Description 1.1, W3C Candidate Recommendation. https://www.w3.org/TR/wot-thing-description/. Accessed: 2026-01-05.

[5] 2026. HomeAssistant home automation platform. https://www.home-assistant.io/. Accessed: 2026-01-05.

[6] 2026. Introduction to openHAB. https://www.openhab.org/docs/. Accesed: 2026-01-05.

[7] 2026. SPADE Agent Development Framework. https://spadeagents.eu/. Accessed: 2026-01-05.

[8] Andrei Ciortea, Olivier Boissier, and Alessandro Ricci. 2019. Engineering world-wide multi-agent systems with hypermedia. In *Engineering Multi-Agent Systems: 6th International Workshop, EMAS 2018, Stockholm, Sweden, July 14-15, 2018, Revised Selected Papers 6*. Springer, 285–301.

[9] Andrei Ciortea, Simon Mayer, Olivier Boissier, and Fabien Gandon. 2019. Exploiting interaction affordances: on engineering autonomous systems for the web of things. (2019).

[10] Hongwei Cui, Yuyang Du, Qun Yang, Yulin Shao, and Soung Chang Liew. 2024. Llmind: Orchestrating ai and iot with llm for complex task execution. *IEEE Communications Magazine* (2024).

[11] Ken Ducatel, Union européenne. Technologies de la société de l'information, Union européenne. Institut d'études de prospectives technologiques, and Union européenne. Société de l'information conviviale. 2001. Scenarios for ambient intelligence in 2010. (2001).

[12] Evan King, Haoxiang Yu, Sangsu Lee, and Christine Julien. 2024. Sasha: creative goal-oriented reasoning in smart homes with large language models. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 8, 1 (2024), 1–38.

[13] Silin Li, Yuhang Guo, Jiashu Yao, Zeming Liu, and Haifeng Wang. 2025. HomeBench: Evaluating LLMs in Smart Homes with Valid and Invalid Instructions Across Single and Multiple Devices. *arXiv preprint arXiv:2505.19628* (2025).

[14] Don Norman. 2013. *The design of everyday things: Revised and expanded edition.* Basic books.

[15] Alessandro Ricci, Michele Piunti, and Mirko Viroli. 2011. Environment programming in multi-agent systems: an artifact-based perspective. *Autonomous Agents and Multi-Agent Systems* 23 (2011), 158–192.

[16] Gyuhyeon Seo, Jungwoo Yang, Junseong Pyo, Nalim Kim, Jonggeun Lee, and Yohan Jo. 2025. Simuhome: A temporal-and environment-aware benchmark for smart home llm agents. *arXiv preprint arXiv:2509.24282* (2025).

[17] Yingtian Shi, Xiaoyi Liu, Chun Yu, Tianao Yang, Cheng Gao, Chen Liang, and Yuanchun Shi. 2024. Bridging the gap between natural user expression with complex automation programming in smart homes. *arXiv preprint arXiv:2408.12687* (2024).

[18] Alexandru Sorici, Andrei Olaru, and Adina Magda Florea. 2025. Towards Agentic AI support for Hypermedia MAS Models of Smart Environments.