A Vision for LLM-based Interaction Assistance for Autonomous Agents on the Semantic Web

Danai Vachtsevanou[1], Jérémy Lemée[1], Valentina Tamma[2], Terry R. Payne[2], Andrei Ciortea[1,3] and Simon Mayer[1]

[1]University of St.Gallen, St. Gallen, Switzerland
[2]The University of Liverpool, Liverpool, UK
[3]Inria, Univ. Côte d'Azur, CNRS, I3S, France

Abstract

Advances in the Semantic Web and the Web of Things have enabled the dynamic advertisement of interaction descriptions, which allow autonomous agents to discover and reason about actions in hypermedia environments. However, these descriptions occasionally fall short in open, dynamic settings—for example, they may contain incomplete knowledge about unexpected situations, or information that does not directly address the needs of heterogeneous agents. In practice, such gaps are typically bridged by humans, often relying on their common sense. In this paper, we envision the integration of Large Language Models (LLMs) into Hypermedia Multi-Agent Systems (MAS) to leverage their world knowledge to fill information gaps at run time and enable scalable support for agent interaction in open and dynamic hypermedia environments. Our proposal lays the groundwork for a framework that considers LLM-based assistive functions for interaction in Hypermedia MAS, enabling the parameterisation and contextual grounding of interaction, methods for agents to leverage this assistance, and safeguards to ensure integrity and transparency. With these ideas, our aim is to inspire further research exploring the extent to which LLMs can assist in managing semantic descriptions in hypermedia environments at run time, while balancing scalable interaction assistance with identified challenges.

Keywords

Semantic Web, Autonomous Agents, Large Language Models, Interaction Assistance

1. Introduction

Research in the Web of Things (WoT) has enabled the semantic representation of possible actions using Semantic Web technologies. These representations allow autonomous agents to discover, reason about, and execute actions in hypermedia environments. They are referred to as *signifiers* [1], to highlight their role as interaction cues—guiding agents by providing easily discoverable and meaningfully interpretable information about available actions. The use and design of signifiers is being explored in Web-based Multi-Agent Systems (MAS), for example through the adoption of W3C WoT Thing Descriptions (TDs) [2], as well as through extensions to TDs or custom approaches adapted for specific agent-processable representation formats, models, reasoning mechanisms, or domains (e.g., [3, 4, 5, 6]).

However, agents still face challenges in effectively and efficiently navigating, observing, and reasoning over a large number of signifiers, whose content may be tailored to different agents or frequently updated to reflect the dynamic action availability and context of physical-virtual settings. To address this, various assistive services and methods have been proposed, e.g., for search, resource and context monitoring,

The Second International Workshop on Hypermedia Multi-Agent Systems (HyperAgents 2025), in conjunction with the 28th European Conference on Artificial Intelligence (ECAI 2025); October 26, 2025, Bologna, Italy

*Corresponding author.
danai.vachtsevanou@unisg.ch (D. Vachtsevanou); jeremy.lemee@unisg.ch (J. Lemée); V.Tamma@liverpool.ac.uk (V. Tamma); T.R.Payne@liverpool.ac.uk (T. R. Payne); andrei.ciortea@unisg.ch (A. Ciortea); simon.mayer@unisg.ch (S. Mayer)
https://www.alexandria.unisg.ch/entities/person/Danai_Vachtsevanou (D. Vachtsevanou); https://www.alexandria.unisg.ch/entities/person/Jeremy_Lemee (J. Lemée); https://www.liverpool.ac.uk/people/valentina-tamma (V. Tamma); https://www.liverpool.ac.uk/people/terry-payne (T. R. Payne); https://andreiciortea.ro/ (A. Ciortea); https://www.alexandria.unisg.ch/entities/person/Simon_Mayer (S. Mayer)
0000-0002-6697-0427 (D. Vachtsevanou); 0000-0003-0828-1188 (J. Lemée); 0000-0002-1320-610X (V. Tamma); 0000-0002-0106-8731 (T. R. Payne); 0000-0003-0721-4135 (A. Ciortea); 0000-0001-6367-3454 (S. Mayer)

filtering, planning, context-aware interaction authorisation, and semantic content negotiation [1, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17]. Many of these approaches rely on deterministic methods (e.g., reasoners, SHACL validation), providing predictable and transparent assistance to agents. However, they may fail in cases where agents encounter unfamiliar vocabularies, evolving ontologies, or incomplete and ambiguous contextual information. In such cases, human designers must intervene, drawing on their domain understanding and common sense to bridge these gaps—for example, by enriching signifiers with additional content or updating them to address new requirements.

Recently, the reported world knowledge of Large Language Models (LLMs) has led researchers to explore their potential in areas such as ontology and knowledge graph engineering [18, 19], Web navigation assistance and search for agents [20, 21, 22], relating actions with objects and effects for robots [23], and recommending actions as part of the reasoning cycle of Web-based agents [24]. In this context, we propose further exploring such functionalities to investigate the extent to which LLMs could assist in guiding agents' actions on the Web. This vision is illustrated through a scenario that highlights challenges agents may face when discovering and reasoning about actions in semantic hypermedia environments—challenges that often require world knowledge to resolve (see Section 2). Building on this, we propose a set of LLM-based functions within a conceptual framework for scalable, LLM-assisted interaction among people and autonomous agents on the Web (see Section 3).

2. Motivating Scenario

To motivate this work, we present a scenario for using devices in an open, dynamic hypermedia environment, which highlights the challenges autonomous agents face in perceiving action possibilities and deciding how to act in such settings. A smart office building uses agents to optimise the working environment based on the ambient conditions (e.g., room lighting conditions, time of day) and the current activities of human inhabitants. One specific activity is to work on a document within a personal space (e.g., on a personal laptop or e-ink reader). Agents can perceive the environment and infer the ongoing activities in the environment. Additionally, the office building provides agents with several actuators, such as those for controlling the lighting and window blinds. In the current context, an agent identifies the goal of changing the room illuminance level to optimise reading conditions. Several representative failure cases (C1–C7) illustrate the potential gaps in semantic, procedural, or contextual knowledge available at run time, which may prevent the agent from achieving its goal.

- C1 *Action Mismatch*: The agent is aware that room illuminance can be increased using a `saref4bldg:Lamp` with the `saref:OnCommand` action. However, the device description exposes the action `schema:ActivateAction`. If the agent does not understand that these two actions are semantically and functionally equivalent, the agent will fail to achieve its goal.
- C2 *Functional Substitutability*: During daylight hours, the agent may increase the illuminance by opening the window blinds (`saref4bldg:ShadingDevice`); for example, using the action `saref:OpenCommand`. However, to do this, the agent needs to recognise alternative actions that can achieve the same goal, and needs to be able to evaluate their suitability in the current context.
- C3 *Goal-Action Mapping*: The agent has adopted the goal of increasing the illuminance in the room, but cannot identify any available relevant action, and thus the goal fails.
- C4 *Contextual Misalignment*: The agent's goal is to facilitate optimal reading conditions given the device used for the activity. If an emissive device is used (e.g., a tablet), then lighting conditions should be low; conversely, they should be greater when using a reflective device (e.g., an e-ink reader). Thus, failing to infer the correct context may result in the wrong action being performed.
- C5 *Device Operation Gap*: The lamp exposes two actions: `saref:OnCommand` for turning on the lamp, and `saref:SetLevelCommand` for adjusting the lamp's brightness. The agent is only aware of the former action that, when used in isolation, fails to provide sufficient illuminance.
- C6 *Action Mediation*: The agent has the procedural knowledge to combine `saref:OnCommand` with `saref:SetLevelCommand` to set a lamp's brightness to a specific level expressed in `dbpedia:Lux`. However, the device expects the brightness input as a `dbpedia:Percentage`. Mediation is needed to convert a `dbpedia:Lux` value into the equivalent `dbpedia:Percentage` for that device.

C7 *Goal Mediation*: The agent implements the Belief-Desire-Intention (BDI) model, and its goal is represented internally through the `increase(illuminance)` predicate. However, the lamp offers a `saref:OnCommand` action that is associated with the goal of increasing illuminance in natural language through an `rdfs:comment`. As the agent cannot interpret natural language, mediation is needed to map this action to its goal.

3. Towards a Framework for LLM-based Interaction Assistance in Hypermedia MAS

We envision a conceptual framework for leveraging LLMs to support agents in perceiving and exploiting actions in Hypermedia MAS, identifying the key *functional roles* of LLMs (e.g., curating Web ontologies and signifiers), and examining the *knowledge sources* needed for contextual grounding, the *integration options* of LLM-based functions in Hypermedia MAS, and the *safeguards* for ensuring semantic integrity and reliability. The core elements explored in each selected topic are captured in Figure 1.

3.1. LLM-based Functions for Interaction Assistance

We propose investigating a set of LLM-based functions as potential enablers to assist agents in addressing interaction challenges, as identified through cases C1-C7 of our motivating scenario (Section 2). These functions could be implemented following a general *template*: a function takes as *input* a set of parameters relevant to the specific problem it is designed to address (e.g., the URIs for `saref:OnCommand` and `schema:ActivateAction` in C1). Using these parameters, the function then *constructs a prompt* that frames the problem for the LLM, typically in the form of a targeted question (for example, *"Do the following semantic annotations refer to the same concept?"*). This prompt may also incorporate domain-specific context, such as Web ontologies or (parts of) knowledge graphs, serialised into a textual format suitable for language model input (e.g., using Turtle [25] or JSON-LD [26]). The *output* of the function is derived from the LLM's response, with optional post-processing applied as necessary.

A function may rely on specific techniques and workflows for prompting LLMs. For example, chain-of-thought prompting allows an LLM to consider its reasoning traces when deriving an output [27]. ReAct enables an LLM to perform sequences of reasoning, actions, and observations before deriving an answer [28]. LLMs can perform actions using tools [29], which may provide additional information to the LLM context [29]. The Model Context Protocol [30] facilitates tool usage and integrating LLMs with data sources. Techniques like Retrieval Augmented Generation (RAG) augment LLMs with an external memory (used during the prompt generation process) to enhance LLMs' responses with factual data [31]. In this context, RAG can be used, for example, to provide additional knowledge (e.g., as Web

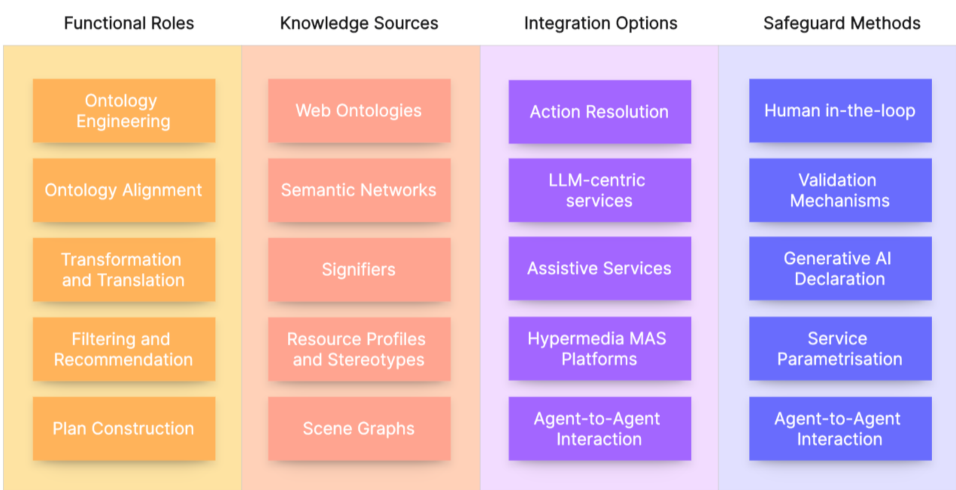

Figure 1: The core elements of our envisioned framework for LLM-assisted interaction in Hypermedia MAS.

ontologies) that is not a function's direct input but retrieved during the generation process. Ontologies could provide natural language descriptions of concepts to help an LLM determine their meanings.

Considering the generic function template and state-of-the-art prompt engineering and orchestration techniques, we examine the functional roles LLMs could play in assisting agents in Hypermedia MAS.

Ontology Construction and Maintenance: Ontology engineering, i.e. the process of designing and maintaining structured representations of knowledge within a domain, is an active research area, resulting in well-established methodologies and approaches [32, 33, 34, 35]. However, designing and constructing such representations is a challenging and time consuming activity [36], and recently, the use of LLMs to support the ontology engineering process has been explored. Specifically, the integration of LLMs with traditional symbolic methods has been investigated, with results suggesting that they can play a key role in knowledge engineering workflows and usher in a new phase of knowledge representation that combines explicit and parametric knowledge [19, 37].

The ontology engineering life-cycle can be broadly split into 4 phases: 1) requirement elicitation and analysis; 2) conceptualization; 3) implementation; and 4) maintenance; and the integration of LLMs has been explored for each of these phases. The requirements analysis and elicitation phase has garnered particular attention, largely due to its persistent susceptibility to the well-known *knowledge acquisition bottleneck* problem [38]. Recent proposals include LLM-based conversational tools that mediate the interactions between ontology engineers and domain experts [39], LLMs supporting text-based approaches [40], the automatic re-engineering of ontology requirements for ontology reuse [41], and extracting requirements from knowledge graphs [42] as well as concept/property–based generation strategies [43]. LLMs have also been exploited to support other activities in the lifecycle, notably concept and axiom definition, and taxonomy enrichment [44, 45], ontology validation [46] and population [47], as well as the whole generation process [48]. Finally, particularly relevant for facilitating interoperability in (Hypermedia) MAS, LLMs have been proposed as a way to align ontologies [49] and to support the development of modular, reusable ontologies [50]. Research on Web-based MAS should thus remain attuned to emerging LLM-based functionalities for key ontology engineering tasks such as concept extraction, axiom generation, and ontology alignment. These could help identify equivalent concepts across ontologies (C1) and goal-action alignments (C4), but also support commonsense reasoning about action [51] through accurate commonsense ontologies, including procedural knowledge (C6).

Curation of Action Relations: These functions could identify and organise relationships between actions within or across signifiers in a Hypermedia MAS. The goal is to inform agents about the meaningful, useful, and valid ordering of actions, that remains up-to-date with features of the environment and the abilities of target agents, as well as the general run-time context that may include: device and service states; agent goals; tasks, etc. For example, a function could take as input the signifiers of a target artifact (e.g., a URI that identifies the lamp description), generate a prompt to elicit appropriate action sequences, and update the relevant signifiers by establishing links between them (e.g., between `saref:OnCommand` and `saref:SetLevelCommand` in C5). Similarly, it could process signifiers of multiple artifacts to infer inter-artifact action dependencies (C6); e.g., linking a signifier of a lamp to a mediator service. For more goal-oriented assistance, the function could optionally accept a goal description to be linked to a relevant signifier, e.g., in C3. LLM-based functions have the potential for curating relations between interaction-related information in additional cases such as relating an action with one that reverses the effects of the former (e.g., `saref:OnCommand` with `saref:OffCommand`), and actions that monitor the execution and effect of another (e.g., `saref:OnCommand` with `saref:GetCommand`).

Contextual Filtering and Recommendation: This class of functions filters irrelevant actions or recommends the most contextually appropriate ones to help agents discover and reason about available options while reducing the agents' workload. This could be achieved by tailoring signifier exposure to the run-time situation and characteristics of both agent and environment. The function could take as input information about a target agent, such as through a URI identifying a profile that includes the agent's goals, actions, procedural knowledge, and related characteristics. In addition, the function should retrieve signifiers from the environment. These different types of information would be used to generate an LLM prompt that requests valid and relevant actions for the agent. Optionally, additional context (such as the current state of artifacts and other agents) could be incorporated to enhance relevance. These functions could support the same cases as an action relation curation function, but without altering the shared signifiers in the hypermedia environment. Instead, they would selectively expose signifiers to target agents. For example, a function could consider the agent's goal (C3), interaction history (C5), and ability set (C6) from the agent's profile to recommend actions such as `saref:OnCommand`, `saref:SetLevelCommand`, and `chameo:DataNormalisation`, respectively.

Translation: These functions translate signifiers from one representation to another to accommodate agents that differ in how they interpret and integrate signifiers to their cognitive processes. Such a function could take as input a signifier to be translated, as well as information about the target representation (e.g., as an identifier or a semantic description of the required format). The LLM could use additional information about the different representation formats considered and how to get from one to the other (e.g., information about how to translate a natural language description to a formal representation in predicate logic if such a representation is possible). For example, this function could assist in C7 by translating the natural language description of the goal associated with an action into the format that the agent uses (e.g., predicate logic), which is provided to the LLM through the agent profile. The agent could then decide whether the provided action can be used to achieve its goal.

3.2. Knowledge Source for Prompting and Contextual Grounding

Access to *core ontologies* is critical across all functional roles, as they provide the foundational vocabulary for representing Hypermedia MAS and available actions. Integrating these into LLM-based functions could enable more appropriate prompt interpretation and modification of signifiers or ontologies themselves. Key ontologies include the W3C WoT TD Ontology [2], which enables defining *signifiers* for properties, actions, and events with their parameters, and the the Hypermedia MAS (hMAS) Core Ontology [52], which provides generic terms for agents, artifacts, workspaces, and organizations. It also defines signifiers, which may offer recommendations regarding the context in which an interaction should be enacted, and the types of agents best suited to perform it. The hMAS ontology may be valuable for handling agent profiles in functions when information about a target agent is required, such as agent goals in filtering and recommendation, or agents' cognitive abilities in translation. It could also support broader contextual grounding by leveraging profiles of other MAS entities. Additional vocabularies could be used, e.g., for action pre- and post-conditions [4], goals [12], and context management [11]. For physical environments, *scene graphs* [53] could offer spatial-temporal representations of objects, their visual (and potentially functional) similarities with other objects, and human intentions [53].

Beyond ontologies and descriptions that reflect the run-time state of a Hypermedia MAS, background knowledge from commonsense resources could support LLM-based functions. General-purpose sources like ConceptNet [54] and ATOMIC [55] offer broad commonsense knowledge, while more specialised ones support physical [56] and social [57] common sense. Commonsense knowledge specifically for interaction assistance could also be explored, e.g., through *artifact stereotypes*, similar to *component stereotypes* [58], capturing behavioral expectations of artifact types based on their intrinsic characteristics and roles in a system. Similarly, *agent stereotypes* could capture expectations for different agent types, including typical goals, behavioral and cognitive abilities linked to specific architectures or roles. Complementary to these, *agent personas* could provide detailed profiles of typical agents in a MAS, including their procedural knowledge and preferences, to ground LLMs' outputs in realistic scenarios.

In general, functions may process both unstructured and structured information, like natural language descriptions, semantic descriptions (e.g., hMAS agent profiles [1]), and JSON documents (e.g., A2A Agent Cards [59]). These may be retrieved at run time by functions, e.g., through input URIs. If ontologies are used to describe such resources, these too could be retrieved by the function to support its processes.

3.3. Integrating LLM-based Assistance in Hypermedia MAS

Different approaches could be employed for integrating LLM-based functions into Hypermedia MAS, with trade-offs for agents and the MAS as a whole. One approach is to integrate LLM-based functionalities in existing agent architectures to support *action resolution*, the process of reasoning about whether an action is relevant and available, by resolving knowledge about abstract actions into executable actions based on discovered signifiers [60]. For example, one architecture was proposed that combines planning with LLMs to generates action recommendations when plans fail [24]. This tight integration ensures full transparency and autonomy over when and how the LLM is invoked for interaction assistance, including complete control of knowledge sources, prompt engineering, and contextual grounding.

LLM-based functions could also be offered as services in the environment that agents discover and use. For example, the LLM-based navigation assistant in [21] is implemented as such a service, taking a Wikipedia page URL and a target keyword, and returning action recommendations to agents navigating Wikipedia. This approach offloads prompt engineering, contextual grounding, and resource management to the service, which can evolve and be optimized independently, and potentially leverage richer data sources and more advanced models. However, such services may not be specialized for the specific needs of different agents, and they reduce control for the agent and its designer.

Additionally, the proposed LLM-based functions could be integrated into existing *assistive services*, such as description directories [61, 62] and search engines [7], or services that offer similar functionalities based on conventional methods, such as for ontology alignment. Extended with LLM-based functions, such services could complement their core methods to more flexibly support cases where conventional methods fall short by delegating to built-in LLM-based functions. For example, this could resolve failures in rule-based Signifier Exposure Mechanisms [12, 1] for action recommendation. Such integration would foster a more seamless agent experience on the Web, with the edge case of LLM-based assistance becoming a default feature of *Hypermedia MAS platforms* (e.g., [10]). The decrease in transparency and control should be mitigated by allowing agents to configure services and access explanations of LLM-driven outcomes. Finally, LLM-based assistant agents should be explored to support through *agent-to-agent interaction*, e.g., through protocols for negotiating ontological correspondances [63].

3.4. Integrity and Transparency Safeguards for LLM-based Interaction Assistance

Several challenges remain closely tied to ongoing research in ontology engineering that relies on hybrid approaches that combine LLMs with traditional symbolic methods. Such approaches require validation to address issues such as the trade-offs between expressivity and decidability [64], rigorously assess techniques that incorporate LLM components [65], and confront challenges related to the limited transparency of LLMs including concerns about reliability and reproducibility [66]. Web ontologies and semantic descriptions generated or curated by LLMs still require human evaluation and integrity verification methods, whether at design time or at run time by dedicated services and capable agents.

A key challenge also lies in ensuring that agents retain visibility and control over the use of LLM-based functions, especially when these are seamlessly embedded within non-LLM-centric services or Hypermedia MAS platforms. This could be addressed through transparent *declarations of Generative AI usage*, aligned with regulations such as the forthcoming Article 50 of the EU AI Act [67], to support agent trust. Moreover, agents should have the option to *configure and parameterise* their interaction with such services and platform, granting them explicit options to select if, when, and to what extent LLM-based assistance is used, potentially exclusively as a fallback mechanism (e.g., as in [24]).

4. Conclusions

In this paper, we explored the potential for LLMs to assist autonomous agents in discovering possible actions and interacting in hypermedia environments, particularly in cases where access to world knowledge would be valuable but its manual provisioning by humans (users or designers) does not scale. Toward this vision, we laid the groundwork for a future framework that should explore functional roles, knowledge requirements, integration options, and safeguards for LLM-assisted interaction in Hypermedia MAS. We argue that research on agents' interactions on the Web should remain closely aligned with advances in LLM-assisted ontology engineering to enhance the use and interpretation of signifiers on the open and dynamic Web. At the same time, it is crucial to develop assistive tools that preserve visibility and control, ensuring these features do not unnecessarily interrupt or complicate agents' interactions, or the programming activities of agent designers.

Declaration on Generative AI

LLMs were used as part of the scientific method of the proposed work, but were not used in the construction or composition of the paper itself.

References

[1] D. Vachtsevanou, A. Ciortea, S. Mayer, J. Lemée, Signifiers as a first-class abstraction in Hypermedia Multi-Agent Systems, in: Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, 2023.

[2] V. Charpenay, M. Kovatsch, S. Käbisch, T. Kamiya, M. McCool, Web of Things (WoT) Thing Description, W3C Recommendation, W3C, 2020. URL: https://www.w3.org/TR/2023/REC-wot-thing-description11-20231205/.

[3] E. O'Neill, D. Lillis, G. M. O'Hare, R. W. Collier, Delivering multi-agent MicroServices using CArtAgO, in: Engineering Multi-Agent Systems: 8th International Workshop, EMAS 2020, Auckland, New Zealand, May 8–9, 2020, Revised Selected Papers 8, 2020.

[4] M. Freund, D. Schraudner, S. Schmid, C. Stade, T. Wehr, A. Harth, The SPA Ontology: Towards a Web of Things ready for robotic agents, in: International Workshop on Actionable Knowledge Representation and Reasoning for Robots 2024, 2024.

[5] H. Ren, D. Anicic, T. A. Runkler, Towards semantic management of on-device applications in industrial iot, ACM Transactions on Internet Technology 22 (2022).

[6] G. Ramanathan, S. Mayer, A match made in semantics: Physics-infused digital twins for smart building automation, in: 2024 IEEE 20th International Conference on Automation Science and Engineering (CASE), 2024.

[7] S. Bienz, A. Ciortea, S. Mayer, F. Gandon, O. Corby, Escaping the streetlight effect: Semantic hypermedia search enhances autonomous behavior in the Web of Things, in: Proceedings of the 9th International Conference on the Internet of Things, 2019.

[8] A. Kamilaris, S. Yumusak, M. I. Ali, WOTS2E: A search engine for a Semantic Web of Things, in: 2016 IEEE 3rd World Forum on Internet of Things (WF-IoT), 2016.

[9] E. O'Neill, R. W. Collier, Exploiting service-discovery and OpenAPI in Multi-Agent MicroServices (MAMS) applications, in: International Workshop on Engineering Multi-Agent Systems, 2023.

[10] A. Ciortea, M. Castellucci, K. Schultz, J. Lemée, D. Vachtsevanou, F. Gandon, S. Mayer, V. Berger, S. Burattini, A. Ricci, O. Boissier, Yggdrasil: An artifact-based framework for hypermedia multi-agent systems, in: R. Collier, V. Mascardi, A. Ricci (Eds.), Agents and Multi-Agent Systems Development – Platforms, Toolkits, Technologies, 2025.

[11] A. Sorici, A. M. Florea, Putting context into hypermedia MAS, in: R. Collier, A. Ricci, V. Nallur, S. Burattini, A. Omicini (Eds.), Multi-Agent Systems, 2025.

[12] J. Lemée, D. Vachtsevanou, S. Mayer, A. Ciortea, Signifiers for conveying and exploiting affordances: From Human-Computer Interaction to Multi-Agent Systems, Annals of Mathematics and Artificial Intelligence 92 (2024).

[13] A. Sorici, A. M. Florea, Towards context-based authorizations for interactions in hypermedia-driven agent environments-the CASHMERE framework, in: Engineering Multi-Agent Systems: 11th International Workshop, EMAS 2023, London, UK, May 29–30, 2023, Revised Selected Papers, 2023.

[14] V. Charpenay, A. Zimmermann, M. Lefrançois, O. Boissier, Hypermedea: A framework for Web (of Things) agents, in: Companion Proceedings of the Web Conference 2022, 2022.

[15] A. Ciortea, S. Mayer, F. Michahelles, Repurposing manufacturing lines on the fly with multi-agent systems for the Web of Things, in: Proceedings of the 17th international conference on autonomous agents and multiagent systems, 2018.

[16] Y. Taghzouti, Semantic content negotiation for knowledge exchange between heterogeneous systems, Ph.D. thesis, Université de Lyon, 2024.

[17] Y. Taghzouti, D. Vachtsevanou, S. Mayer, A. Ciortea, A step toward semantic content negotiation., in: 23rd International Conference on Knowledge Engineering and Knowledge Management (Companion), 2022.

[18] D. Garijo, M. Poveda-Villalón, E. Amador-Domínguez, Z. Wang, R. García-Castro, Ó. Corcho, LLMs for ontology engineering: A landscape of tasks and benchmarking challenges, in: R. Alharbi, J. de Berardinis, P. Groth, A. Meroño-Peñuela, E. Simperl, V. Tamma (Eds.), Proceedings of the Special Session on Harmonising Generative AI and Semantic Web Technologies (HGAIS 2024) co-located with the 23rd International Semantic Web Conference (ISWC 2024), Baltimore, Maryland, November 13, 2024, volume 3953 of *CEUR Workshop Proceedings*, 2024.

[19] B. P. Allen, F. Ilievski, S. Joshi, Identifying and consolidating knowledge engineering requirements, 2023. URL: https://arxiv.org/abs/2306.15124. arXiv:2306.15124.

[20] R. Nakano, J. Hilton, S. Balaji, J. Wu, L. Ouyang, C. Kim, C. Hesse, S. Jain, V. Kosaraju, W. Saunders, X. Jiang, K. Cobbe, T. Eloundou, G. Krueger, K. Button, M. Knight, B. Chess, J. Schulman, WebGPT: Browser-assisted question-answering with human feedback, CoRR abs/2112.09332 (2021). URL: https://arxiv.org/abs/2112.09332. arXiv:2112.09332.

[21] D. Vachtsevanou, J. Lemée, R. Rot, S. Mayer, A. Ciortea, G. Ramanathan, Hyperbrain: Human-inspired hypermedia guidance using a large language model, in: Proceedings of the 34th ACM Conference on Hypertext and Social Media, HT '23, 2023.

[22] A. PBC, How we built our multi-agent research system, 2025. URL: https://www.anthropic.com/engineering/built-multi-agent-research-system.

[23] F. Ocker, J. Deigmöller, J. Eggert, Exploring Large Language Models as a Source of Common-Sense Knowledge for Robots, arXiv e-prints (2023). doi:10.48550/arXiv.2311.08412. arXiv:2311.08412.

[24] S. Schmid, M. Freund, A. Harth, Adaptive planning on the web: Using LLMs and affordances for web agents, in: International Knowledge Graph and Semantic Web Conference, Springer, 2024.

[25] G. Carothers, E. Prud'hommeaux, RDF 1.1 Turtle, W3C Recommendation, W3C, 2014. URL: https://www.w3.org/TR/2014/REC-turtle-20140225/.

[26] P.-A. Champin, D. Longley, G. Kellogg, JSON-LD 1.1, W3C Recommendation, W3C, 2020. Https://www.w3.org/TR/2020/REC-json-ld11-20200716/.

[27] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou, et al., Chain-of-thought prompting elicits reasoning in large language models, Advances in neural information processing systems 35 (2022).

[28] S. Yao, J. Zhao, D. Yu, N. Du, I. Shafran, K. Narasimhan, Y. Cao, React: Synergizing reasoning and acting in language models, in: International Conference on Learning Representations (ICLR), 2023.

[29] G. Mialon, R. Dessì, M. Lomeli, C. Nalmpantis, R. Pasunuru, R. Raileanu, B. Rozière, T. Schick, J. Dwivedi-Yu, A. Celikyilmaz, et al., Augmented language models: a survey, arXiv preprint arXiv:2302.07842 (2023).

[30] Model Context Protocol, Specification, Anthropic PBC, 2025. URL: https://modelcontextprotocol.io/specification/2025-06-18.

[31] P. Lewis, E. Perez, A. Piktus, F. Petroni, V. Karpukhin, N. Goyal, H. Küttler, M. Lewis, W.-t. Yih, T. Rocktäschel, et al., Retrieval-augmented generation for knowledge-intensive nlp tasks, Advances in neural information processing systems 33 (2020).

[32] M. Fernández-López, A. Gómez-Pérez, N. Juristo, Methontology: From ontological art towards ontological engineering, in: Proceedings Of the AAAI Spring Symposium on Ontological Engineering, AAAI97, 1997.

[33] N. Noy, et al., Ontology development 101: A guide to creating your first ontology, Technical Report, Stanford knowledge systems laboratory technical report KSL-01-05, 2001.

[34] M. C. Suárez-Figueroa, A. Gómez-Pérez, M. Fernández-López, The NeOn methodology framework: A scenario-based methodology for ontology development, Applied ontology 10 (2015).

[35] V. Presutti, E. Daga, A. Gangemi, E. Blomqvist, Extreme design with content ontology design patterns, in: Proc. of the 2009 International Conf. on Ontology Patterns, volume 516, 2009.

[36] J. F. Sequeda, W. J. Briggs, D. P. Miranker, W. P. Heideman, A pay-as-you-go methodology to design and build enterprise knowledge graphs from relational databases, in: Proceedings of the 18th International Semantic Web Conference, ISWC 2019, 2019.
[37] B. AlKhamissi, M. Li, A. Celikyilmaz, M. T. Diab, M. Ghazvininejad, A review on language models as knowledge bases, 2022. URL: https://arxiv.org/abs/2204.06031. arXiv:2204.06031.
[38] E. A. Feigenbaum, Knowledge engineering: The applied side of artificial intelligence, Annals of the New York Academy of Sciences 426 (1984).
[39] B. Zhang, et al., Ontochat: A framework for conversational ontology engineering using language models, in: Proc. of the 21st Extended Semantic Web conference, ESWC, 2025.
[40] M. Antia, C. M. Keet, Automating the generation of competency questions for ontologies with agocqs, in: Knowledge Graphs and Semantic Web, 2023.
[41] R. Alharbi, et al., An experiment in retrofitting competency questions for existing ontologies, in: Proc. of the 39th ACM/SIGAPP Symposium on Applied Computing, SAC '24, 2024.
[42] F. Ciroku, et al., Revont: Reverse engineering of competency questions from knowledge graphs via language models, Journal of Web Semantics 82 (2024).
[43] Y. Rebboud, et al., Can LLMs generate competency questions?, in: Proc. of the 21st Extended Semantic Web conference, ESWC, 2024.
[44] M. Funk, S. Hosemann, J. C. Jung, C. Lutz, Towards ontology construction with language models, in: S. Razniewski, J. Kalo, S. Singhania, J. Z. Pan (Eds.), Joint proceedings of the 1st workshop on Knowledge Base Construction from Pre-Trained Language Models (KBC-LM) and the 2nd challenge on Language Models for Knowledge Base Construction (LM-KBC) co-located with the 22nd International Semantic Web Conference (ISWC 2023), Athens, Greece, November 6, 2023, volume 3577 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2023. URL: https://ceur-ws.org/Vol-3577/paper16.pdf.
[45] H. Dong, J. Chen, Y. He, Y. Gao, I. Horrocks, A language model based framework for new concept placement in ontologies, in: A. Meroño-Peñuela, A. Dimou, R. Troncy, O. Hartig, M. Acosta, M. Alam, H. Paulheim, P. Lisena (Eds.), The Semantic Web - 21st International Conference, ESWC 2024, Hersonissos, Crete, Greece, May 26-30, 2024, Proceedings, Part I, volume 14664 of *Lecture Notes in Computer Science*, 2024. URL: https://doi.org/10.1007/978-3-031-60626-7_5. doi:10.1007/978-3-031-60626-7_5.
[46] S. Tsaneva, S. Vasic, M. Sabou, Llm-driven ontology evaluation: Verifying ontology restrictions with chatgpt, in: S. Tiwari, N. Mihindukulasooriya, F. Osborne, D. Kontokostas, J. D'Souza, M. Kejriwal, M. A. Pellegrino, A. Rula, J. E. L. Gayo, M. Cochez, M. Alam (Eds.), Joint proceedings of the 3rd International workshop on knowledge graph generation from text (TEXT2KG) and Data Quality meets Machine Learning and Knowledge Graphs (DQMLKG) co-located with the Extended Semantic Web Conference (ESWC 2024), Hersonissos, Greece, May 26-30, 2024, volume 3747 of *CEUR Workshop Proceedings*, 2024.
[47] G. Ciatto, A. Agiollo, M. Magnini, A. Omicini, Large language models as oracles for instantiating ontologies with domain-specific knowledge, Knowledge-Based Systems 310 (2025).
[48] A. S. Lippolis, M. J. Saeedizade, R. Keskisärkkä, S. Zuppiroli, M. Ceriani, A. Gangemi, E. Blomqvist, A. G. Nuzzolese, Ontology generation using large language models, in: European Semantic Web Conference, 2025.
[49] H. B. Giglou, J. D'Souza, F. Engel, S. Auer, Llms4om: Matching ontologies with large language models, CoRR abs/2404.10317 (2024). URL: https://doi.org/10.48550/arXiv.2404.10317. doi:10.48550/ARXIV.2404.10317. arXiv:2404.10317.
[50] C. Shimizu, P. Hitzler, Accelerating knowledge graph and ontology engineering with large language models, Journal of Web Semantics 85 (2025).
[51] R. T. Icarte, J. A. Baier, C. Ruz, A. Soto, How a general-purpose commonsense ontology can improve performance of learning-based image retrieval, in: Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17, 2017.
[52] A. Ciortea, F. Gandon, M. Lefrançois, D. Vachtsevanou, C. Leturc, O. Boissier, L. G. Nardin, S. Mayer, V. Charpenay, A. Zimmermann, Hypermedia MAS Core Ontology, Technical Report, HyperAgents,

2021. URL: https://purl.org/hmas/core.
[53] S. S. Jha, K. Garcia, Y. S. Antille, M. E. Solèr, S. Padua, S. Mayer, Temporal scene understanding using contextually unique identification, in: 2024 IEEE 36th International Conference on Tools with Artificial Intelligence (ICTAI), 2024.
[54] R. Speer, J. Chin, C. Havasi, Conceptnet 5.5: An open multilingual graph of general knowledge, in: Proceedings of the AAAI conference on artificial intelligence, volume 31, 2017.
[55] J. D. Hwang, C. Bhagavatula, R. Le Bras, J. Da, K. Sakaguchi, A. Bosselut, Y. Choi, (comet-) atomic 2020: On symbolic and neural commonsense knowledge graphs, in: Proceedings of the AAAI conference on artificial intelligence, volume 35, 2021.
[56] Y. Bisk, R. Zellers, J. Gao, Y. Choi, et al., Piqa: Reasoning about physical commonsense in natural language, in: Proceedings of the AAAI conference on artificial intelligence, volume 34, 2020.
[57] M. Sap, H. Rashkin, D. Chen, R. Le Bras, Y. Choi, Social IQa: Commonsense reasoning about social interactions, in: K. Inui, J. Jiang, V. Ng, X. Wan (Eds.), Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), Association for Computational Linguistics, Hong Kong, China, 2019.
[58] G. Ramanathan, S. Mayer, Reasoning about physical processes in buildings through component stereotypes, in: Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation, BuildSys '23, Association for Computing Machinery, New York, NY, USA, 2023. URL: https://doi.org/10.1145/3600100.3623737. doi:10.1145/3600100.3623737.
[59] Agent2Agent (A2A) Protocol Specification v0.2.5, Specification, Google LLC, 2025. URL: https://a2aproject.github.io/A2A/v0.2.5/specification/.
[60] D. Vachtsevanou, B. de Lima, A. Ciortea, J. F. Hübner, S. Mayer, J. Lemée, Enabling BDI agents to reason on a dynamic action repertoire in hypermedia environments, in: Proceedings of the 2024 International Conference on Autonomous Agents and Multiagent Systems, International Foundation for Autonomous Agents and Multiagent Systems, 2024. doi:10.5555/3635637.3663048.
[61] K. Toumura, M. McCool, A. Cimmino, F. Tavakolizadeh, Web of Things (WoT) Discovery, W3C Recommendation, W3C, 2023. URL: https://www.w3.org/TR/2023/REC-wot-discovery-20231205/.
[62] F. for Intelligent Physical Agents (FIPA), FIPA Agent Management Specification, FIPA specification, FIPA, 2000. URL: http://www.fipa.org/specs/fipa00023/SC00023J.html.
[63] T. R. Payne, V. Tamma, Negotiating over ontological correspondences with asymmetric and incomplete knowledge, in: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems, 2014.
[64] J. Z. Pan, S. Razniewski, J.-C. Kalo, S. Singhania, J. Chen, S. Dietze, et al., Large language models and knowledge graphs: Opportunities and challenges, Transactions on Graph Data and Knowledge 1 (2023).
[65] B. P. Allen, L. Stork, P. Groth, Knowledge engineering using large language models, Transactions on Graph Data and Knowledge 1 (2023).
[66] E. La Malfa, A. Petrov, S. Frieder, C. Weinhuber, R. Burnell, R. Nazar, et al., Language models as a service: Overview of a new paradigm and its challenges, 2023. URL: https://arxiv.org/abs/2309.16573. arXiv:2309.16573.
[67] Article 50: Transparency obligations for providers and deployers of certain ai systems, Artificial Intelligence Act (Regulation (EU) 2024/1689), Official Journal version of 13 June 2024 (2024). URL: https://artificialintelligenceact.eu/article/50/.