Converging Hypermedia, Protocols, and Knowledge Architectures: A New Paradigm for Grounded and Interoperable LLM-Agent Systems

Maksim Ilin[1]

[1]Transport and Telecommunication Institute, Engineering Faculty, Riga, Latvia

Abstract

The evolution of LLM-powered agents has led to a distinction between single-purpose AI Agents and collaborative Agentic AI systems. While powerful, both face persistent challenges in grounding, interoperability, and governance when operating in open Web environments. This paper argues that a new, more robust paradigm is emerging from a deep convergence of architectural principles and modern standards. We posit that Hypermedia Multi-Agent Systems (HMAS) provide the foundational philosophy for Web-native agency, where the concepts of W3C Web of Things (WoT) "Things" and Model Context Protocol (MCP) "Tools" are unified as complementary specializations of hypermedia-described Web resources with discoverable affordances. This unified view enables the design of advanced Agentic AI systems—composed of autonomous AI Agents grounded in foundational definitions of agency—that can reliably perceive and act upon a rich environmental fabric of physical (WoT) and semantic (DKG) resources. By analyzing this convergence and presenting an illustrative agentic architecture, we demonstrate how this paradigm offers enhanced grounding, true interoperability, and clearer pathways to trustworthy governance, providing a structured approach for developing the next generation of scalable and context-aware autonomous systems.

Keywords

Hypermedia Multi-Agent Systems, Model Context Protocol, LLM Agents, Web of Things, Distributed Knowledge Graphs, Agent Architectures, Interoperability, Grounding, Semantic Web

1. Introduction

The proliferation of powerful Large Language Models (LLMs) has spurred the development of "Agentic AI", where LLMs serve as the cognitive core for autonomous agents. Frameworks such as AutoGen, CrewAI, and Magentic-One [1] demonstrate this potential by enabling multi-agent teams to tackle complex tasks through orchestrated collaboration. However, as these systems aspire to operate beyond predefined teams and engage with the dynamic, open complexity of the World Wide Web, significant challenges persist. These include ensuring robust factual grounding for LLM outputs, achieving true interoperability across diverse systems and data silos, and establishing trustworthy governance and operational safety [2]. Many current agent orchestration platforms, while adept at managing internal team dynamics, often lack inherent mechanisms for dynamic discovery and interaction within truly open, decentralized environments, a core tenet of Web-native agency.

The Hypermedia Multi-Agent Systems (HMAS) paradigm, championed by the HyperAgents community [3], offers a visionary blueprint for agents that perceive, reason, and act by navigating a Web structured as a discoverable hypermedia environment. This paper posits that a new, more potent and practically realizable paradigm for autonomous agents, especially LLM-based Agents, is emerging from the synergistic convergence of:

- Foundational HMAS principles promoting dynamic, decentralized interaction via hypermedia affordances.
- Modern agent interaction and data access protocols, with the Model Context Protocol [4] serving as a key enabler for standardized agent-tool communication.

ilin.m@tsi.lv (M. Ilin)
https://tsi.lv (M. Ilin)
0009-0006-1955-4777 (M. Ilin)

- Advanced Web architectures, specifically the Web of Things [5] for grounding agents in physical reality, and Distributed Knowledge Graphs (DKGs) – enhanced with hybrid vector-graph capabilities [6] and techniques like Denoising Diffused Embeddings [7] and Interlocked Hypergraphs [8] – for providing rich semantic context and verifiable knowledge.

This paper argues that the path forward lies not in viewing these components as separate pillars to be integrated, but in recognizing a deeper convergence. We propose a unified paradigm where HMAS provides the foundational philosophy for Web-native agency, and where WoT "Things" and MCP "Tools" are understood as complementary specializations of a single concept: the hypermedia-described Web resource. We detail this unified view, define our terminology for agents and services based on foundational literature, and present an illustrative Agentic AI system architecture. We demonstrate how this converged paradigm offers inherent mechanisms for enhanced LLM grounding, interoperability, and governance. Finally, we identify key research directions, aiming to contribute to the workshop's objective of building conceptual and technological bridges between the AI and Web communities.

2. A Unified Paradigm for Web-Native Agency

The proposed paradigm is best understood as a unified architectural vision grounded in the principles of the Web itself. This vision is defined by its core philosophy (HMAS), its unified view of resources (Things and Tools), and the rich environments these resources inhabit.

2.1. Defining Agents, Services, and Systems

To clarify our discussion, we adopt a precise taxonomy. We distinguish a simple program from an autonomous agent based on the foundational definition by Franklin and Graesser [9], which requires an agent to be situated, sense and act on its environment over time in pursuit of its own agenda, and affect its future sensing. Building on this, we adopt the modern, LLM-era taxonomy of Sapkota et al. [10]:

- **Service or Tool:** A reactive program that exposes capabilities (affordances) via an interface (e.g., an MCP Server, a WoT device API). It lacks its own overarching agenda and temporal continuity.
- **AI Agent:** A single, autonomous, tool-using entity that satisfies the Franklin & Graesser criteria, often with an LLM as its cognitive core.
- **Agentic AI System:** A collaborative multi-agent system, composed of specialized AI Agents, working under an orchestration layer to achieve complex, high-level goals.

2.2. HMAS: A Foundational Philosophy for Web-Native Agency

The HMAS paradigm [3] provides the architectural philosophy. It envisions agents as first-class citizens of the Web, operating within a distributed hypermedia environment. Key principles include Environmental Interaction, HATEOAS (Hypermedia as the Engine of Application State), and Dynamic Discovery of resources and their affordances, allowing agents to "escape the streetlight effect" of only interacting with pre-known entities [11].

2.3. A Unified View of WoT Things and MCP Tools

A core insight is that WoT "Things" and MCP "Tools" are not fundamentally different categories; they are complementary specializations of the same underlying concept: a *Hypermedia-Described Web Resource with Discoverable Affordances.* As Table 1 illustrates, both define interaction models, description formats, and protocols for agents to perceive and act. Their primary divergence lies in their target domains and non-functional requirements: WoT is optimized for resources tied to the physical world, while MCP is optimized for interaction with LLM agents, emphasizing rich context management. A prime example of this convergence is a WoT-enabled robotic lab controlled via a standardized MCP interface, where a physical "Thing" exposes its capabilities as a "Tool".

Table 1
Comparison of W3C WoT Things and MCP Tools as Hypermedia Resources.

Aspect	WoT Things	MCP Tools
Core Definition	Abstraction of a physical/virtual entity with metadata and interfaces.	Standardized interface for an AI model to access an external tool or data source.
Interaction Model	Properties (state), Actions (functions), Events (notifications).	Discovery (list tools), Execution (call tool), real-time updates.
Description Format	JSON-LD Thing Description with rich semantic vocabulary.	JSON-RPC 2.0 with JSON Schema for tool name, description, and inputs.
Semantic Layer	Deep integration with RDF, ontologies, and Linked Data.	Schema-based descriptions without inherent Semantic Web integration.
Target Domain	IoT devices, cyber-physical systems.	AI applications, LLM integration, digital services.

2.4. Rich Web Environments: Grounding in Physical and Semantic Reality

This unified view of resources is grounded in rich environments that provide context and data:

- **Web of Things (WoT):** Standardized by the W3C, WoT extends the hypermedia fabric to the physical world, allowing agents to discover, understand, and interact with sensors and actuators in a standardized way [12]. Agents can even reason about incomplete WoT descriptions using existential reasoning [5].
- **Distributed Knowledge Graphs (DKGs):** Built on Semantic Web technologies, DKGs serve as the semantic backbone, providing structured knowledge. Modern DKGs leverage Hybrid Vector-Graph Databases [6] and are enhanced with techniques like DDE [7] and Interlocked Hypergraphs [8] for completeness and can be made verifiable using blockchain [13].

3. Illustrative Architecture: An Agentic Knowledge Retrieval System

To illustrate the converged paradigm, we conceptualize an "Agentic Knowledge Retrieval System," inspired by author's ongoing development efforts like Knowledge Retrieval Agentic System and OpenAlex Explorer MCP Server projects. This system is composed of an ensemble of specialized AI Agents, each satisfying the criteria for autonomous agency as defined by Franklin and Graesser [9]. This distinguishes them from the reactive services they utilize, such as the "OpenAlex MCP Server". The architecture, shown in Figure 1, can be implemented using frameworks such as CrewAI and LangGraph.
Key Components & Workflow:

1. **Orchestration Layer** manages the overall process, delegating tasks to the specialized AI Agents.
2. **External Knowledge Access (MCP):** The SearchAgent, a specialized AI Agent, receives a research query. It exhibits autonomy in executing its search plan and is temporally continuous for the duration of its task. It acts by invoking the "OpenAlex MCP Server" service via MCP.
3. **Knowledge Ingestion & Representation (DKG & Vector DB):** The IngestionAgent, upon receiving data, autonomously executes its goal of processing and storing it. It extracts metadata into a Neo4j DKG and text embeddings into a Qdrant Vector DB.
4. **Grounded Synthesis (LLM & RAG):** The SynthesisAgent, a cognitive AI Agent with an LLM core, pursues its goal of generating a response. It exhibits autonomy by first retrieving context via RAG from the DKG and Vector DB, and then synthesizing its grounded answer.
    - Querying the Neo4j DKG for structured information (e.g., influential authors, citation patterns).
    - Performing semantic similarity searches in Qdrant for relevant text passages. (Access to Neo4j/Qdrant further could be itself realized via internal MCP tools).

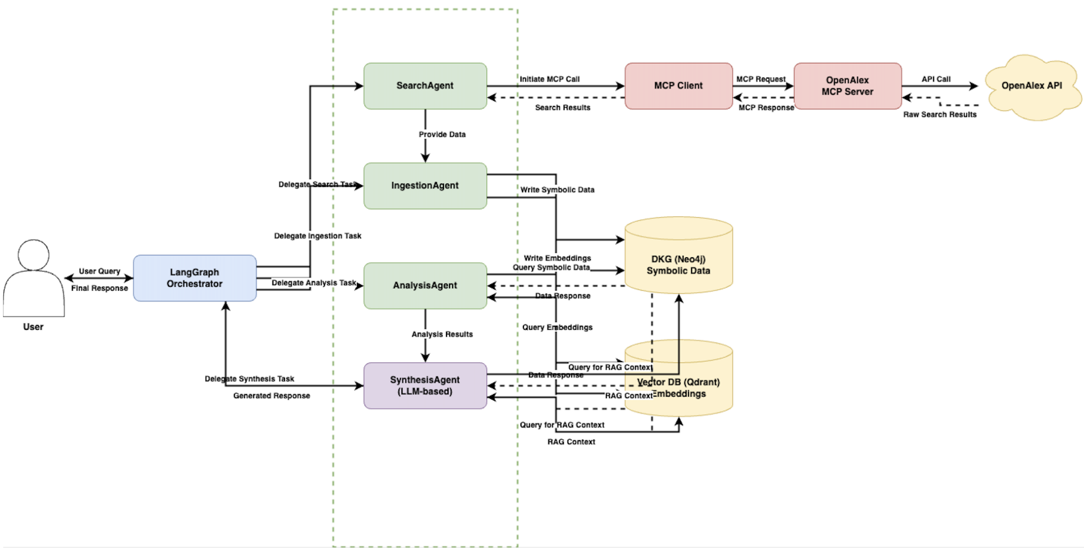

Figure 1: Conceptual Architecture of a Converged Agentic Knowledge Retrieval System.

5. **State Management:** Each AI Agent maintains its own task-specific state, while the orchestrator manages the overall workflow state, demonstrating a multi-agent system with social ability through coordination.

This system, by leveraging standardized protocols for external interactions and rich internal knowledge structures, exemplifies how agents can perform complex, knowledge-intensive tasks with improved grounding and modularity.

4. Key Benefits & Addressed Challenges of the Convergence

The proposed converged paradigm offers substantial benefits, directly addressing persistent challenges in developing advanced autonomous agents:

- **Significantly Enhanced LLM Grounding:** This is arguably the most critical benefit. LLM-based Agents often struggle with factual accuracy and "hallucinations" [14, 15]. The converged paradigm provides multiple grounding layers: real-time data from WoT, structured factual knowledge and complex relationships from DKGs (further enriched by DDE for hypergraph sparsity) and Interlocked Hypergraphs (for cross-domain knowledge fusion), and verifiable interactions with specialized tools via Model Context Protocol (MCP). This rich, multi-faceted context drastically improves the reliability and trustworthiness of LLM agent outputs.
- **Improved Interoperability and Dynamic Coordination:** Current agent ecosystems are often fragmented. The convergence fosters interoperability at multiple levels:
  - HMAS principles encourage discoverable services and common Web interaction patterns.
  - MCP standardizes agent-tool interfaces.
  - Protocols like Agent-to-Agent (A2A) aim to standardize inter-agent dialogues [16].
  - WoT standardizes device interaction.
  - DKGs with shared ontologies provide semantic interoperability. This multi-layered approach helps mitigate protocol bridging challenges, enabling flexible, dynamic coordination.
- **Pathways to Robust Trust, Security, and Governance:**
  - The transparency of HMAS interactions with structured MCP logs, enables comprehensive hypermedia-enabled audit trails, supporting Explainable Agent Design (XAI) [17].

– Secure protocols (MCP with ETDI Framework [18], A2A with DIDs) and verifiable data layers (DKGs with blockchain-anchored provenance like OriginTrail DKG, or coordination secured by BlockAgents [19]) address trust and security concerns.
– This provides a stronger foundation for effective AI governance and building user trust.

5. Conclusion & Future Directions

The convergence of Hypermedia MAS principles with modern interaction protocols (like MCP) and advanced Web architectures (WoT, enhanced DKGs) represents a significant evolutionary step towards more capable, grounded, interoperable, and trustworthy autonomous Web agents. This paradigm shift, illustrated conceptually by our agentic knowledge retrieval architecture, offers a structured approach to leveraging the strengths of LLMs while mitigating their weaknesses through deep integration with verifiable external knowledge and standardized interaction mechanisms.

This vision's realization depends on progress in several key areas. Future research directions include:

1. **Comprehensive Standardization:** Accelerate efforts within bodies like the W3C WebAgents Community Group to develop and promote standards for HMAS affordance descriptions, MCP-A2A/ANP interoperability profiles, semantic alignment for WoT-DKG integration, and federated DKG protocols.
2. **Scalable and Secure Converged Architectures:** Research into scalable, secure implementations of all components, including robust protocol bridging fabrics, federated DKG solutions (e.g., building on FedMSGL [20]), and zero-trust models spanning the entire converged stack.
3. **Deep WoT Integration for Physical Agency:** Expand research into practical applications of WoT within HMAS-MCP architectures, enabling agents to robustly perceive, reason about, and act upon the physical world, including addressing real-time safety and overall energy efficiency.
4. **Advanced HAI and Governance for Converged Systems:** Develop intuitive HAI (Human-AI Interaction) methods and comprehensive governance frameworks (including stateful monitoring and XAI) tailored to the complexities of these converged, LLM-driven agent ecosystems.

Addressing these directions will pave the way for a new generation of Web agents capable of realizing many aspects of the original Semantic Web vision [21], truly transforming the Web into a collaborative space for both humans and intelligent machines.

Declaration on Generative AI

During the preparation of this work, author used several generative AI tools. To support the initial research phase, Perplexity AI and Google Gemini were employed for Generate literature review and Content enhancement by helping to identify and synthesize concepts from source literature for the author's review. To validate the conceptual framework, tools including GitHub Copilot and Qoder provided Content enhancement by assisting with code generation and quality improvement (linting) for software prototypes. Additionally, Google Gemini 2.5 Pro was used to Paraphrase and reword source material into private research notes. For the manuscript itself, Apple Writing assisted with the author's original text by performing Grammar and spelling check and offering suggestions to Improve writing style and Paraphrase and reword for clarity. The author confirms that no manuscript text was drafted or generated by AI and takes full responsibility for the originality and conclusions of this publication.

References

[1] S. Joshi, Review of autonomous systems and collaborative AI agent frameworks (2025).
[2] M. A. Ferrag, N. Tihanyi, M. Debbah, From LLM reasoning to autonomous AI agents: A comprehensive review, arXiv preprint arXiv:2504.19678 (2025).

[3] A. Ciortea, S. Mayer, F. Gandon, O. Boissier, A. Ricci, A. Zimmermann, A Decade in Hindsight: The Missing Bridge Between Multi-Agent Systems and the World Wide Web, in: Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '19), 2019, p. 5. Blue Sky Ideas Track.
[4] Anthropic, Model Context Protocol Documentation, https://modelcontextprotocol.io, 2024. Accessed: 2025-01-01.
[5] V. Charpenay, S. Käbisch, H. Kosch, Characterizing Web of Things Interactions with Existential Reasoning, Semantic Web ((n.d.)).
[6] S. Liu, Z. Zeng, L. Chen, A. Ainihaer, A. Ramasami, S. Chen, Y. Xu, M. Wu, J. Wang, TigerVector: Supporting Vector Search in Graph Databases for Advanced RAGs, arXiv preprint arXiv:2501.11216 (2025).
[7] S. Wu, J. Yang, G. Xu, J. Zhu, Denoising Diffused Embeddings: a Generative Approach for Hypergraphs, arXiv preprint arXiv:2501.01541 (2025).
[8] K. Boughanmi, A. Ansari, Y. Li, EXPRESS: Modeling Categorized Consumer Collections with Interlocked Hypergraph Neural Networks, Journal of Marketing Research ((n.d.)).
[9] S. Franklin, A. Graesser, Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents, in: International workshop on agent theories, architectures, and languages, Springer, 1996, pp. 21–35.
[10] R. Sapkota, K. I. Roumeliotis, M. Karkee, Ai agents vs. agentic ai: A conceptual taxonomy, applications and challenges, arXiv preprint arXiv:2505.10468 (2025).
[11] 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, pp. 1–8. doi:10.1145/3365871.3365901.
[12] J. A. Martins, A. Mazayev, N. Correia, Hypermedia APIS for the Web of Things, IEEE Access 5 (2017) 20058–20067. doi:10.1109/ACCESS.2017.2755259.
[13] M. Pawlak, M. Stolarczyk, A. Poniszewska-Marańda, Food Supply Chain Management with Blockchain Technology in Implementation of Hyperledger Fabric, in: 2025 IEEE 22nd International Conference on Software Architecture Companion (ICSA-C), 2025, pp. 261–270.
[14] C. Agostinho, Explainability as the Key Ingredient for AI Adoption in Industry 5.0 Settings, Frontiers in Artificial Intelligence 6 (2023). doi:10.3389/frai.2023.1264372.
[15] M. R. Smeets, P. G. Roetzel, R. J. Ostendorf, AI and its Opportunities for Decision-Making in Organizations: A Systematic Review of the Influencing Factors on the Intention to use AI, Die Unternehmung 75 (2021) 432–460. doi:10.5771/0042-059X-2021-3-432.
[16] A. Ehtesham, A. Singh, G. K. Gupta, S. Kumar, A survey of agent interoperability protocols: Model context protocol (mcp), agent communication protocol (acp), agent-to-agent protocol (a2a), and agent network protocol (anp), arXiv preprint arXiv:2505.02279 (2025).
[17] S. Rodriguez, J. Thangarajah, A. Davey, Design patterns for explainable agents (XAg), in: Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems, 2024, pp. 1621–1629.
[18] V. S. Narajala, I. Habler, Enterprise-grade security for the model context protocol (MCP): Frameworks and mitigation strategies, arXiv preprint arXiv:2504.08623 (2025).
[19] B. Chen, G. Li, X. Lin, Z. Wang, J. Li, BlockAgents: Towards Byzantine-Robust LLM-Based Multi-Agent Coordination via Blockchain, in: Proceedings of the ACM Turing Award Celebration Conference-China 2024, 2024, pp. 187–192. doi:10.1145/3674399.3674445.
[20] D. Li, Z. Yang, S. Xie, FedMSGL: A Self-Expressive Hypergraph Based Federated Multi-View Learning, in: Proceedings of the AAAI Conference on Artificial Intelligence, 2025, pp. 18244–18252.
[21] T. Berners-Lee, J. Hendler, O. Lassila, The Semantic Web, in: Linking the World's Information: Essays on Tim Berners-Lee's Invention of the World Wide Web, 2023, pp. 91–103. Reprint of original 2001 article.