
Position: Collaborative Agentic AI Needs Interoperability Across Ecosystems

Anonymous Authors¹

Abstract

Collaborative agentic AI is projected to transform entire industries by enabling AI-powered agents to autonomously perceive, plan, and act within digital environments. Yet, current solutions in this field are all built in isolation, and we are rapidly heading toward a landscape of fragmented, incompatible ecosystems. In this position paper, we argue that **interoperability, achieved by the adoption of minimal standards, is essential to ensure open, secure, web-scale, and widely-adopted agentic ecosystems**. To this end, we devise a minimal architectural foundation for collaborative agentic AI, named WEB OF AGENTS, which is composed of four building blocks: agent-to-agent messaging, interaction interoperability, state management, and agent discovery. WEB OF AGENTS adopts existing standards and reuses existing infrastructure where possible. With WEB OF AGENTS, we take a first but critical step toward interoperable agentic systems and offer a pragmatic path forward before ecosystem fragmentation becomes the norm.

1. Introduction

Agentic artificial intelligence (AI) is an emerging approach in which autonomous agents, typically powered by large language models (LLMs), can make decisions, take actions, and learn from the environment in which they are deployed (Acharya et al., 2025). The capabilities of such agents go beyond generative AI as they can invoke tools, maintain long- and short-term memory, and engage in complex workflows with other agents (Liu et al., 2025; Murugesan, 2025). Their ability to autonomously plan, communicate, and collaborate unlocks a new class of possibilities ranging from scientific discovery to multi-agent enterprise orchestration (Ghareeb et al., 2025; Li et al., 2023; Lu et al.,

2024; Park et al., 2023). As academia and industry starts to understand these systems better, the impact of agentic AI is poised to be transformative.

While individual agents have shown tremendous potential with the use of tools such as deep research (OpenAI, 2025) and web search (Perplexity, 2025), multi-agent AI systems have recently gained traction to enable *collaborative agentic AI* (Liu et al., 2025; Sapkota et al., 2025). Recent deployments demonstrate their potential across scientific domains: AI Scientist (Lu et al., 2024) coordinates specialized agents for end-to-end research automation, ROBIN (Ghareeb et al., 2025) integrates planning and experimental agents for biomedical discovery, and comparable systems accelerate materials discovery through coordinated simulation and modeling agents (Bazgir et al., 2025). As these agentic capabilities proliferate across domains, the next frontier is a decentralized WEB OF AGENTS where agents can dynamically discover each other’s expertise and establish communication to tackle complex, cross-disciplinary challenges that no single agent could address alone (Malika et al., 2024; Shlezinger et al., 2021).

Currently, there is a significant effort to build infrastructure for collaborative agentic AI. Notable examples include the Agent2Agent (A2A) protocol (Google DeepMind, 2025), released by Google in April 2025, the Model context protocol (MCP) (Model Context Protocol Project, 2025), released by Anthropic in November 2024, and the Agent communication protocol (ACP) (Agent Communication Protocol Project, 2025), released by IBM in April 2025 (merged into A2A later). These solutions introduce necessary primitives such as agent-to-agent communication protocols, agent authentication methods, and mechanisms for expressing the capabilities of agents. However, these siloed efforts develop in isolation with distinct abstractions and data formats, creating a fragmented landscape where agents from one ecosystem cannot interoperate or interact with agents from another. We visualize this challenge in Figure 1 (left). While such diverse and decentralized efforts in implementation accelerates innovation and experimentation (Eiras et al., 2024; Raymond, 1999), our concern is that the resulting ecosystem fragmentation imposes engineering overheads, limits reusability, and threatens long-term scalability and security.

Two obvious approaches could address this fragmentation:

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.




Figure 1. Left: Prominent protocols for collaborative agentic AI (A2A (Google DeepMind, 2025), ANP (Agent Network Protocol Project, 2025) and ACP (Agent Communication Protocol Project, 2025)). Agents by default cannot communicate with agents in other ecosystems, resulting in fragmentation (see Section 3). Right: our proposed solution, WEB OF AGENTS.

(i) enforcing a single unified protocol, or (ii) creating post-hoc translation layers between each pair of systems. However, history suggests that both approaches are problematic. The former stifles innovation and often fails to gain adoption, as seen with the OSI model’s struggle against TCP/IP’s pragmatic success (Clark, 1988; Leiner et al., 1997; 2009; Shafer, 2014), while the latter creates undesirable complexity and maintenance burdens (Berners-Lee, 1999; Bollinger, 2000; Chen et al., 2023; Kolos, 2000; Zhao et al., 2023). In this position paper, we argue that the field of collaborative agentic AI needs a third approach to address fragmentation. We strongly advocate for a WEB OF AGENTS: an interoperable ecosystem of collaborative agents that are grounded in minimal standards. We visualize this vision in Figure 1 (right). More concretely, **our position is that interoperability, achieved by the adoption of minimal standards, is essential to ensure open, secure, web-scale, and widely-adopted agentic ecosystems.** Rather than constraining development by enforcing a particular protocol or creating brittle translation layers, we advocate for the adoption of lightweight standards that already power web infrastructure. This interoperability enables several key benefits: reduced duplication of effort, scalable ecosystem growth, enhanced security through shared scrutiny, and greater autonomy for users and agent operators as they can freely switch between agents that provide similar services.

To realize this vision, we identify four minimal, yet sufficient building blocks that ecosystems should adopt for enabling interoperability: (i) agent-to-agent messaging, enabling asynchronous and structured communication via existing web protocols, (ii) interaction interoperability, establishing shared interaction specifications through interaction documents, (iii) state management, supporting both short-term sessions and long-term state, and (iv) discovery, allowing agents to find one another. Our approach requires minimal new infrastructure, just careful extensions of the web as we know it. Together, these building blocks form the foundation of the WEB OF AGENTS.

Contributions. To defend our position, we first analyze several prominent solutions in Section 3, identifying common trends across current frameworks for collaborative agentic AI. This analysis motivates our arguments in Section 4 that

interoperability is not merely desirable, but essential for realizing the full potential of collaborative agentic AI. We then propose in Section 5 a minimal blueprint, WEB OF AGENTS, that demonstrates how this interoperability can be achieved through careful reuse and extension of existing web technologies. We open-source our implementation of applying the standards to A2A, ACP, and Agora, including a demo¹. Finally, we discuss some views that oppose our position in Section 6 before concluding with final remarks in Section 7.

2. From Generative AI to Collaborative Agentic AI

Generative AI. Generative AI encompasses computational models capable of producing novel content across multiple modalities, including text, images, and audio (Feuerriegel et al., 2024). Among the most prominent implementations is ChatGPT, which provides a conversational interface to large language models (LLMs) through natural language processing (Achiam et al., 2023). LLMs demonstrate proficiency in diverse language generation tasks, including document summarization, content synthesis, and creative text generation. The standard generative AI workflow, illustrated in Figure 2 (left), follows a request-response pattern where users submit queries to an LLM system (Step 1) and receive generated outputs (Step 2). This architecture supports multi-turn conversations, enabling contextual dialogue through sequential query-response exchanges within a persistent session.

However, this conventional user-to-LLM paradigm exhibits notable architectural limitations (Schneider, 2025). A fundamental constraint is the models’ inability to perform autonomous environmental observation, external system interaction, or action execution without explicit human orchestration. Such capabilities would encompass database querying, web navigation (Drouin et al., 2024), or programmatic transaction execution. While retrieval-augmented generation (RAG) frameworks (Lewis et al., 2020) partially address these limitations by augmenting model inputs with relevant external information, these approaches typically require separate orchestration mechanisms to manage the retrieval process.

Agentic AI. To address these architectural constraints, the field has witnessed a paradigm shift toward agentic artificial intelligence systems (Hughes et al., 2025; Murugesan, 2025). Agentic AI architectures are fundamentally structured around autonomous *agents*—computational entities capable of independent operation and dynamic interaction with external *tools*, as illustrated in Figure 2 (middle). These agents are typically specialized and optimized for specific

¹Source code available at https://anonymous.4open.science/r/web_of_agents_demo-B011/

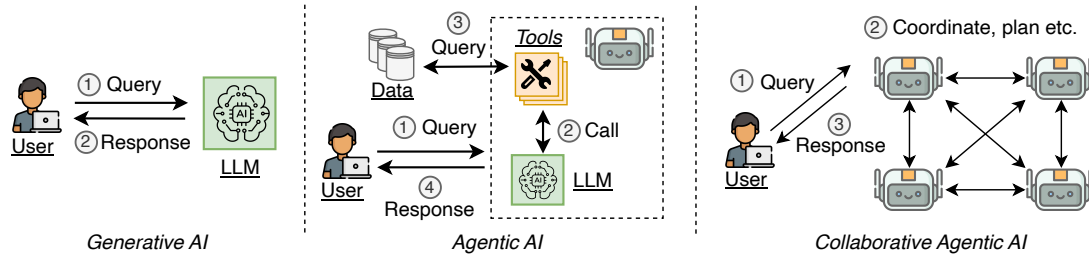


Figure 2. Generative AI (left), agentic AI (middle), and collaborative agentic AI (right). This work provides a blueprint for interoperable collaborative agentic AI that leverages existing web protocols.

task domains, such as computer vision processing or web navigation protocols. Tools constitute external service interfaces that agents can programmatically invoke to execute functions beyond the inherent capabilities of the underlying LLMs. Tool selection and orchestration represent critical components of the response generation pipeline, requiring systematic evaluation of available resources for optimal task completion (Ruan et al., 2023). Agents can autonomously invoke multiple tools in parallel (Step 2 in Figure 2) when required, utilizing tool responses to inform subsequent planning and action selection. These tools maintain direct interfaces to various data repositories and external systems (Step 3). Agentic AI frameworks demonstrate significant potential for transformative applications across domains including healthcare and software engineering (Murugesan, 2025).

However, contemporary agentic architectures predominantly operate through single-agent paradigms, constraining system scalability, modularity, and fault tolerance. This architectural limitation necessitates that individual agents manage comprehensive task execution, logical reasoning, and tool orchestration, potentially exceeding agent capabilities when confronting complex scenarios requiring specialized domain expertise or access to heterogeneous tool ecosystems.

Collaborative Agentic AI. Collaborative agentic AI represents the next step and involves AI agents that can communicate and collaborate across tasks, domains, and organizational boundaries. We visualize this frontier in Figure 2 (right). Compared to agentic AI with a single agent, agents can form collectives and work together (Step 2). For example, a travel agent can coordinate with a calendar agent to check availability, a payment agent to process transactions, and a policy agent to enforce organizational constraints. These interactions can all occur without human intervention. Such networks of LLM-based agents have already been used to simulate human and economic behavior (Li et al., 2023; Park et al., 2023), and to understand dynamics in societies (Gürcan, 2024). The recent boom in AI agents follows substantial work on multi-agent systems (MAS), a field with decades of research on autonomous agents that, similar to collaborative agentic AI, coordinate, negotiate,

and act in digital environments (Ferber & Weiss, 1999; Finin et al., 1994; FIPA, 1997; 2002). Classical MAS research focuses on pre-programmed agents with static rules, whereas LLM-powered agents can operate in more unstructured environments and reason using natural language. Therefore, while we acknowledge the field of classical MAS, we position this work towards LLM-powered agents.

3. The Alarming Trend Toward Incompatible Agentic Ecosystems

The potential of (collaborative) agentic AI has resulted in significant research and engineering efforts from both industry and academia to devise new infrastructure, protocols, and frameworks (Hosseini & Seilani, 2025; Tran et al., 2025). Notable examples include Google’s A2A protocol (Google DeepMind, 2025), Anthropic’s MCP (Model Context Protocol Project, 2025), and IBM’s ACP (Agent Communication Protocol Project, 2025), leading to the “protocol wars” phase of (collaborative) agentic AI. Based on existing work (Ehtesham et al., 2025) and our own research, we identify prominent solutions for (collaborative) agentic AI and look at their innovation areas, see Table 1. Even though most solutions focus on inter-agent communication, we include MCP and agents.json in our analysis because of their increased adoption in agentic workflows (Ehtesham et al., 2025). The goal of our analysis is not to point out particular shortcomings of individual solutions but rather to take a broader perspective on the development of agentic AI solutions.

Our main observation is that **agentic AI solutions are developed in isolation and are incompatible with each other.** As a result, we are heading toward a collection of ecosystems in which agents built using one framework are unlikely to interoperate with those from another (Research, 2024). Such fragmentation introduces inefficiencies, for example, companies wishing to deploy agents that interact across systems must either commit to a single protocol or reimplement agents for each distinct ecosystem. As new solutions continue to emerge, these interoperability challenges are poised to grow.

One might address the interoperability challenges among

diverse agentic AI systems with *translation layers* or *mediators* that convert messages and actions between different protocols (Paniagua, 2021). This approach, for example, has been adopted to connect fragmented web standards in the early days of the web and to enable value exchange between different blockchain platforms. While this brings interoperability, the viability of translational layers has been questioned (Berners-Lee, 1999; Kolos, 2000). Firstly, translation layers can become brittle or outdated as protocols evolve. Secondly, such layers can expand the system’s attack surface and expose vulnerabilities that malicious actors could exploit, as has been observed with cross-blockchain bridges (Zhao et al., 2023). Thirdly, as the number of distinct protocols increases, maintaining translation mappings becomes increasingly difficult (Chen et al., 2023). These challenges raise serious concerns about whether translation layers can serve as a long-term solution for agentic AI interoperability.

Beyond interoperability challenges, this ecosystem-level isolation has also led to uneven progress across the agentic AI stack. A holistic solution for agentic AI requires, among others, primitives such as agent identity, protocols for inter-agent communication, agent discovery mechanisms, and mechanisms for memory management. We observe in Table 1 discrepancies in the focus of innovation areas across solutions. On the one hand, all identified solutions emphasize some notion of agent identity, such as proposing mechanisms for agents to describe their capabilities or roles. On the other hand, only a minority of solutions implement mechanisms for the discovery of other agents in the network. Similarly, the capability of an agent to manage short- or long-term memory receives inconsistent attention across solutions. This uneven focus results in fragmented tooling that hinders composability across systems. This pattern is not unique to agentic AI: similar fragmentation has occurred in domains such as internet-of-things (Aly et al., 2018), blockchain (Zamyatin et al., 2021), and early web services (Papazoglou & Van Den Heuvel, 2007), where numerous tools emerged to solve narrow problems in isolation.

Some analyzed solutions introduce their own implementations of core primitives such as authentication, capability exposure, agent discovery, and message exchange formats. For example, while some leverage existing authentication schemes such as the OAuth protocol (Hardt, 2012) or decentralized identifiers (DIDs) (Reed et al., 2020), others either rely on custom mechanisms or leave this aspect unspecified. Similarly, agent discovery is often realized through custom listings with their own formats. Generally, the result is a patchwork of ecosystems that are incompatible by design. Moreover, some protocols introduce new primitives without clearly defined use cases, optimizing for hypothetical scenarios rather than concrete needs. LMOS, for example, follows a maximalist design by building a complete OS-

like environment for LLMs. Messaging applications are a pertinent analogy to this observation: major services like iMessage, WhatsApp, Signal, or Telegram all implement their own protocol stack, message format, and identity systems, and are largely incompatible with each other (Arnold & Schneider, 2017).

The field of agentic AI stands at a critical point. In the absence of coordination, it risks splintering into competing silos, each solving similar problems in slightly different and incompatible ways. For instance, consider a travel-booking agent developed by a major airline that supports rich capabilities such as real-time seat selection using a particular agent protocol. Even if this agent is exceptionally capable, users and other agents relying on a different agentic ecosystem would be unable to invoke it or compose it into larger workflows. In this way, incompatibility directly burdens engineers and diminishes the quality of service for end users.

Key Takeaway

Because each solution is being developed in isolation, agentic AI is evolving into a landscape of incompatible ecosystems.

4. Interoperability as Necessary Foundation for Agentic AI

If fragmentation is the current trajectory, interoperability is the antidote. To achieve this, we advocate for the adoption of **minimal standards** that enable agents to interact across different ecosystems. Minimal standardization has historically enabled breakthrough innovations across domains: protocols like HTTP and HTML bootstrapped the web (Berners-Lee, 1999), SMTP and IMAP enabled global email interoperability (Crispin, 2003; Klensin, 2008), and TCP/IP continues to serve as the backbone of the Internet (Postel, 1981). With minimal, interoperable foundations, we strongly believe we can bootstrap a successful agentic ecosystem. Next, we further motivate interoperability and elaborate on how it strengthens agentic AI in terms of security, open participation, and scalability.

Security. Interoperability through minimal standardization contributes to the security of the overall ecosystem (Elkhodr et al., 2016). This standardization creates opportunities for systematic auditing, formal verification, and community-wide threat modeling. In fact, security is a critical concern in the emerging landscape of agentic AI and is a necessity to protect against adversarial agents or users (Shavit et al., 2023). This is because, unlike traditional systems that follow predefined logic, AI agents operate in dynamic environments, interpret ambiguous instructions,

Table 1. Prominent solutions for (collaborative) agentic AI and their innovation areas.

Solution	Agent Identity	Inter-Agent Comm.	Agent Discovery	State Management	Interoperability w/ other solutions
A2A (Google DeepMind, 2025)	✓AgentCard	✓Async.	✦Static	✗	✗
MCP (Model Context Protocol Project, 2025)	✦Contextual	✗	✗	✗	✗
ACP (Agent Communication Protocol Project, 2025)	✓	✓REST API	✗	✗	✗
agents.json (Wildcard AI, 2025)	✦Capabilities	✗	✗	✗	✗
Agora (Marro et al., 2024)	✦Protocol docs	✓	✗	✗	✗
ANP (Agent Network Protocol Project, 2025)	✓Identity Layer	✓	✓	✗	✗
AITP (NEAR AI, 2025)	✓	✓Chat Threads	✗	✓	✗
LMOS (Eclipse Foundation, 2025)	✓	✓	✓	✓	✗

and make autonomous decisions, often without human supervision. This autonomy introduces novel attack surfaces, e.g., prompt injection and jailbreak attacks (Liu et al., 2024; Wei et al., 2023). As these agents become embedded in critical workflows, security becomes paramount.

Open Participation. Interoperability empowers agent builders, e.g., developers, organizations, or infrastructure providers, to exercise meaningful choice in the agents they interact with (Kapoor et al., 2025). In a collaborative agentic AI system, agents often depend on other agents. With shared protocols and minimal standards in place, builders can freely choose among competing agents that offer similar functionality but differ in trust models, performance, cost, or affiliation. Without such freedom of choice, switching costs rise dramatically as each agent may require custom APIs, data formats, and authentication schemes. This not only leads to vendor lock-in, but also stifles innovation by discouraging new entrants from building compatible alternatives (Yan & Feng, 2023). This also holds for end users, who may want to directly use the services of agents across ecosystems or combine agents from different providers into a single workflow. An interoperable agentic AI ecosystem fosters open participation from both engineers and end-users, enabling vibrant competition, lowering barriers to entry, and promoting long-term ecosystem health (Aly et al., 2018).

Closely related, interoperability is also a prerequisite for a decentralized agentic AI ecosystem. Without a shared set of standards, coordination inevitably gravitates toward a small number of centralized platforms that dictate how agents are built and deployed. In contrast, minimal interoperability standards allow diverse actors to independently develop and operate agents, while still participating in the overall ecosystem (Kapoor et al., 2025). With interoperability, anyone can publish and run agents that are discoverable and usable by others without needing permission from any authority.

Scalability. Interoperability is also a key enabler of scalability in terms of ecosystem expansion. As the agentic AI landscape grows, no single provider will be able to build and maintain every agent, tool, or integration needed to support diverse user needs across domains. A scalable ecosystem




Figure 3. Blueprint of WEB OF AGENTS and its building blocks.

depends on the ability of independent developers to contribute agents, services, and infrastructure components that can plug into a larger system without needing to coordinate closely with every other participant. Shared protocols reduce the integration overhead, allowing developers to focus on functionality rather than glue code. This mirrors the success of the web, where interoperable standards enabled a vast ecosystem of independent websites, browsers, and services to emerge without centralized planning (Berners-Lee, 1999). Furthermore, it allows the protocols to grow with the necessary improvements as the field evolves and particular use-cases become more prominent.

Key Takeaway

Interoperability is not just a nice-to-have feature, it is foundational. It enables secure, composable, scalable, and decentralized agentic ecosystems across domains.

5. Architectural Foundations for an Interoperable Web of Agents

We now identify a minimal set of building blocks and devise a blueprint for an interoperable agentic AI ecosystem named WEB OF AGENTS. The WEB OF AGENTS architecture is visualized in Figure 3, and the building blocks are summarized in Table 2. This blueprint adheres to the long-standing system design principle: *Keep It Simple, Stupid* (Lampson, 1983). This principle guides our strategy for enforcing *minimal standardization*, and *adopting existing standards* where possible. We strongly believe that the greatest interoperability can be achieved by staying as close as possible to existing web technologies. To reinforce our arguments, we

275 discuss the proof-of-concept implementation for WEB OF
276 AGENTS in Appendix B.

277 5.1. Agent-to-Agent Messaging

278 Standardized message exchange protocols are essential for
279 interoperability in an agentic AI ecosystem, as they establish
280 the communication channels for coordination, negotiation,
281 and collaboration. Without common messaging standards,
282 each pair of interacting agents would require custom inte-
283 gration, creating a scaling problem that would severely
284 restrict ecosystem growth (Papazoglou & Van Den Heuvel,
285 2007).

286
287
288 **Leveraging Hypertext Transfer Protocol (HTTP) for**
289 **agent-to-agent communication.** HTTP, the protocol that
290 underpins the modern web, is well-suited for agent-to-
291 agent communication through its request-response pattern
292 and well-defined semantics (Fielding & Reschke, 2014).
293 By adopting HTTP as the foundational transport protocol,
294 agents can exchange information using established meth-
295 ods: GET requests to retrieve information from other agents
296 (similar to browsers accessing webpages), and POST re-
297 quests for more complex interactions with messages and
298 other context data in the payload. This HTTP-based ap-
299 proach leverages decades of rethinking, improvements, and
300 optimizations without duplication of efforts (Mogul, 2002).
301 While some emerging protocols like ACP and Agora di-
302 rectly leverage HTTP requests for agent-to-agent communi-
303 cation, other solutions like A2A have opted for more special-
304 ized protocols like JSON-RPC (Google DeepMind, 2025;
305 Marro et al., 2024). We strongly believe that agent-to-agent
306 communication does not require specialized protocols like
307 JSON-RPC or WebSockets. While these protocols offer
308 some compelling features such as asynchronous communi-
309 cation, enforcing these as a standard would be an overkill,
310 hinder progress, and limit interoperability (Murley et al.,
311 2021).

312 Another advantage of utilizing HTTP requests for agentic
313 communication is that agents can coexist with webpages,
314 making it trivial to integrate agentic AI systems with exist-
315 ing software. Furthermore, specialized patterns like asyn-
316 chronous requests can be easily implemented on the client-
317 side using HTTP requests and have long existed on the
318 web (Microsoft, 2025). Finally, by building upon HTTP’s
319 extension mechanisms, this foundation allows for future
320 enhancements without breaking backward compatibility,
321 creating a stable yet evolvable communication layer for the
322 WEB OF AGENTS.

323 5.2. Interaction Interoperability

324
325
326 Communication between autonomous agents presents
327 unique challenges beyond basic message transportation.
328 Agents are currently being developed using diverse frame-
329

works, architectural approaches, underlying technologies,
and with varying capabilities (see Section 3). In fact, differ-
ent agents might need completely different inputs, possibly
in different formats, to carry out their tasks. For example, a
weather forecasting agent may expect location coordinates
and a timestamp, while a financial analysis agent requires
historical market data and risk parameters. Existing solu-
tions typically do not sufficiently consider differences in
agent design and capabilities. Protocols typically assume
a universal interaction pattern without accounting for vari-
ability in inputs, expected conditions, or optimal message
compositions that can be efficiently interpreted by the under-
lying LLM (Agent Communication Protocol Project, 2025;
Google DeepMind, 2025). Therefore, collaborative agen-
tic AI requires a way to enable interoperability in agent
interactions.

Interaction document. To enable interaction interoperabil-
ity, we propose an *interaction document* exchange between
agents. This way, each agent can explicitly document its
interface requirements, capabilities, expected inputs, out-
puts, and conditions necessary for effective communication.
Crucially, this documentation need not follow a schema
or predefined structure, where we differ from the classical
MAS works such as KQML and FIPA (Finin et al., 1994;
FIPA, 1997). Since modern AI agents are powered by LLMs,
they can interpret and adapt to various documentation for-
mats, from structured JSON to natural language descriptions.
This documentation acts as a standardized handshake, ac-
cessible via common endpoints using HTTP requests, and
enables agents to understand and adapt to diverse interaction
patterns without prior agreements or custom engineering
efforts by the agent operator. This approach is minimalistic
in the sense that we neither restrict how agents interact with
each other nor impose rigid formatting requirements on their
capability descriptions.

This minimalistic approach to interaction interoperability
offers three advantages. First, it enables agents to evolve
independently, *i.e.*, the internal implementations can change
completely as long as the documentation interface remains
stable. Second, it fosters innovation by avoiding prema-
ture and overly restrictive standards, providing agent op-
erators flexibility in their implementations. Ultimately, this
straightforward yet effective solution mitigates the friction
associated with agent integration.

330 5.3. State Management

331 Agent interactions are inherently stateful, distinguishing
332 them from traditional stateless web requests (Ghosh, 2025).
333 Agents engage in conversations where each interaction
334 changes the internal state of the agent or the environment.
335 This requires both short-term retention of state (*e.g.*, for each
336 pairwise agent interaction) and long-term state persistence.

Table 2. WEB OF AGENTS building blocks.

Building blocks	Functional needs	Web technologies
Agent-to-agent messaging	HTTP-based messaging	HTTP requests
Interaction interoperability	Interaction documentation	API documentation
State management	Short-term memory Long-term memory	Sessions DB integration
Agent discovery	Unique endpoints Capability advertisement	URLs, DNS Well-known paths

However, most agentic AI frameworks adopt stateless designs or delegate state management entirely to the developers (Agent Communication Protocol Project, 2025; Google DeepMind, 2025; Marro et al., 2024). Without standardized approaches to state management, agents would remain confined to isolated and context-free exchanges, severely limiting their interoperability and potential to collaborate on complex tasks (Russell & Norvig, 2016).

Short-term memory through sessions. Agents require mechanisms to maintain immediate context across multiple message exchanges in a pairwise conversation. This requirement can be satisfied through existing web session management patterns, which are well-established. More specifically, short-term state in web interactions is managed through cookies, which is a standardized approach as part of the HTTP protocol, where a unique session identifier is assigned during the initial interaction and subsequently included in all requests within the same conversation (Peng & Cisna, 2000).

Long-term memory through databases. Beyond immediate conversational context, agents often require long-term state storage to store historical interactions and accumulated knowledge across sessions (Sapkota et al., 2025). This long-term memory requirement can be satisfied through established patterns for database integration in web services. The agents can write to and read from the database as part of their workflow. Similar patterns underpin virtually all data-persistent web applications, from social media platforms to collaborative document editing systems.

Thus, we do not need sophisticated or new solutions to state management in agentic AI. Reusing existing approaches for state management offers advantages over prescriptive standards. It allows agent developers to be sure that all agents will have state management, but leaves them the flexibility of implementing it according to their requirements.

5.4. Discovery

Discovery is a foundational requirement for any interoperable agent ecosystem. Without standardized discovery mechanisms, agent ecosystems would fragment into isolated islands of functionality, severely limiting their collective capabilities. Current methods either involve developing

isolated networks of agents like Naptha AI (Naptha AI, 2025) or curating a centralized registry of existing AI agents like Agency Marketplace (Agency, 2025). While the former requires all agents to be on the same closed network, the latter requires manual registry self-curation. Both these approaches are not scalable, and hence, unsuitable for WEB OF AGENTS. To enable a truly collaborative WEB OF AGENTS, we need search engines for agents that regularly crawl and index agents on the web. This mirrors how web search engines transformed the early internet from a collection of disconnected pages into a navigable and discoverable information space (Brin & Page, 1998; Schwartz, 1998).

Unique endpoints. Agents require globally unique identifiers that can be reliably resolved to network endpoints for other agents or search engines to discover and interact with them. Uniform resource locators (URLs) provide an ideal solution as they are already designed to uniquely identify resources on the web. This approach benefits from the established domain name system (DNS) infrastructure for name resolution without introducing any new identification scheme. This approach provides several advantages: (i) URLs are human-readable, (ii) they are hierarchical, allowing organizational namespacing of agents, (iii) the entire URL resolution infrastructure, including the DNS, has been optimized and maintained to handle the scale of the web, and (iv) this design allows agents and webpages to coexist over the web, enabling seamless interoperability.

Capability advertisement. Agents must expose metadata about their capabilities and available tools. Most agentic protocols, including A2A, ACP, and agents.json, describe their formal standards for the description of agents and their capabilities. However, the number of fields in such a formal description can range from basic capabilities to guardrails, and is only expected to grow as we build new agentic AI systems (Casper et al., 2025). Therefore, keeping it simple, each agent should expose a short document that describes its capabilities and tools. This document can potentially be the same as the interaction documentation described in Section 5.2. Similar to A2A, we propose hosting the capability description at a well-known path (similar to *robots.txt* and *sitemap.xml*). This standardized location ensures that both search engines and other agents can reliably locate capability information.

By adopting these design principles for agent discovery, we establish a minimalistic standardization framework that directly leverages and builds on top of the web’s highly optimized and scalable infrastructure. Each component can keep evolving to make WEB OF AGENTS better while maintaining interoperability across agentic AI ecosystems.

5.5. Putting it Together

The four building blocks of WEB OF AGENTS form a sufficient foundation for collaborative agentic AI. To illustrate this, consider the AI Scientist MAS (Lu et al., 2024), where specialized AI agents collaborate for scientific discovery. In WEB OF AGENTS, these agents would operate without prior knowledge of other agents and their capabilities. The discovery mechanism regularly indexes all agents. When the hypotheses-generation agent (H) needs computational support, it can query the discovery to locate a suitable code-generation agent (C) for collaboration. Agent-to-agent messaging enables H to send requests to C , while interaction interoperability allows H to understand C ’s interface requirements and compose properly formatted requests. Finally, state management allows H and C to maintain contextual continuity across multi-turn dialogs without retransmitting complete conversations. Critically, both H and C operate independently and may reside in different agentic ecosystems (e.g., A2A and ACP), yet these four building blocks are minimal and sufficient to enable seamless collaboration. The current nascent state of agentic AI frameworks is a perfect moment for adopting these minimal standards to avoid the *protocol wars* and lay the foundations of an interoperable WEB OF AGENTS.

6. Alternative Views

Alternative View 1: *Web technologies were designed for content delivery and should be reinvented for agentic AI.*

We reject the assumption that web technologies are fundamentally incompatible with agentic AI. In fact, web technologies have evolved from serving static websites to dynamic content. At their core, agents are like dynamic websites: they expose some capabilities, respond to structured requests, maintain internal state, and can be discovered and invoked using standard web protocols. Some agentic AI systems are already extending mature web technologies, although in isolation. Thus, agentic AI requires alignment with web technologies, which we have shown is feasible in Section 5.

Alternative View 2: *It is too soon to worry about interoperability in Agentic AI: natural selection will eventually determine the best practices and standards.*

While it is true that the agentic AI space is still evolving and diverse designs as well as use cases are expected, de-

ferring interoperability until a mature protocol emerges is shortsighted. History shows that dominant solutions often arise without prioritizing openness (e.g., early cloud APIs or messaging platforms). Google’s A2A may become influential, but it is not inherently neutral nor guaranteed to support open collaboration across providers. Messaging apps serve as a cautionary tale: while dominant platforms emerged, the lack of interoperability led to closed ecosystems, lock-in, and duplicated engineering effort (Yan & Feng, 2023). We thus ask ourselves the question: do we want to monopolize agentic AI by adopting one ecosystem, or allow a decentralized existence of diverse but interoperable ecosystems? Our position takes the decentralized stance: early alignment around minimal standards ensures flexibility and guards against future silos, even if popular solutions emerge later.

Alternative View 3: *Standardizing early slows down progress.*

We agree that standardization can slow innovation, especially in a fast-moving field like agentic AI. We acknowledge this and hence, do not advocate for extensive protocols, but for minimal and extensible foundations, just enough to enable interoperability without stifling experimentation. WEB OF AGENTS is the first step towards this. The web itself evolved this way: early standards like HTTP and HTML were minimalistic and were later extended with the uses that were common on the web.

7. Final Remarks and Call to Action

In this work, we argue that collaborative agentic AI demands interoperability as a foundational requirement. Rather than waiting for dominant solutions to emerge, we advocate for the adoption of minimal standards that enable agents to discover, communicate, and collaborate across organizational and technical boundaries (see Appendix C). Our proposed WEB OF AGENTS blueprint shows that such interoperability can be achieved by reusing existing web technologies.

Ongoing efforts by AGNTCY, W3C, and The Linux Foundation demonstrate community momentum toward interoperability (see Appendix A for details). To realize an interoperable agentic AI ecosystem, we identify concrete steps for key stakeholders. *Protocol Developers* must align protocols with the four WEB OF AGENTS building blocks, i.e., adopt HTTP as the transport layer, standardize interaction documentation endpoints, implement web session management, and expose capabilities via URLs. *Standards Organizations* like W3C and The Linux Foundation must prioritize these minimal standards as baseline requirements for hosted protocols. Coordinated action is needed now, before fragmentation becomes the norm. By adopting these minimal standards early, we can build the foundations for an open collaborative agentic AI ecosystem.

References

- 440
441
442 Acharya, D. B., Kuppan, K., and Divya, B. Agentic ai:
443 Autonomous intelligence for complex goals—a compre-
444 hensive survey. *IEEE Access*, 2025.
- 445 Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I.,
446 Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S.,
447 Anadkat, S., et al. Gpt-4 technical report. *arXiv preprint*
448 *arXiv:2303.08774*, 2023.
- 449
450 Agency. Agency marketplace. <https://marketplace.agency.com/agents/>,
451 2025. Accessed: 2025-05-21.
- 452
453 Agent Communication Protocol Project. Agent
454 communication protocol specification. [https://agentcommunicationprotocol.dev/](https://agentcommunicationprotocol.dev/introduction/welcome)
455 [introduction/welcome](https://agentcommunicationprotocol.dev/introduction/welcome), 2025. Accessed:
456 2025-05-20.
- 457
458 Agent Network Protocol Project. Agent net-
459 work protocol (anp) specification. <https://agent-network-protocol.com/>,
460 2025. Accessed: 2025-05-20.
- 461
462 Agency.org. Building infrastructure for the Internet of
463 Agents. <https://agntcy.org/>, 2025. Accessed:
464 2026-01-29.
- 465
466 Aly, M., Khomh, F., Guéhéneuc, Y.-G., Washizaki, H., and
467 Yacout, S. Is fragmentation a threat to the success of the
468 internet of things? *IEEE Internet of Things Journal*, 6(1):
469 472–487, 2018.
- 470
471 Arnold, R. and Schneider, A. An app for every step: A
472 psychological perspective on interoperability of mobile
473 messenger apps. 2017.
- 474
475 Bazgir, A., chandra Praneeth Madugula, R., and Zhang, Y.
476 Multicrossmodal automated agent for integrating diverse
477 materials science data, 2025. URL <https://arxiv.org/abs/2505.15132>.
- 478
479 Berners-Lee, T. *Weaving the Web: The original design and*
480 *ultimate destiny of the World Wide Web by its inventor*.
481 Harper San Francisco, 1999.
- 482
483 Bollinger, T. A guide to understanding emerging interoper-
484 ability technologies. Technical report, 2000.
- 485
486 Brin, S. and Page, L. The anatomy of a large-scale hyper-
487 textual web search engine. *Computer networks and ISDN*
488 *systems*, 30(1-7):107–117, 1998.
- 489
490 Casper, S., Bailey, L., Hunter, R., Ezell, C., Cabalé,
491 E., Gerovitch, M., Slocum, S., Wei, K., Jurkovic, N.,
492 Khan, A., et al. The ai agent index. *arXiv preprint*
493 *arXiv:2502.01635*, 2025.
- 494
Chen, Y., Annebicque, D., Philippot, A., Carré-Ménétrier,
V., and Daneau, T. Evaluation methodology of interoper-
ability for the industrial domain: Standardization vs.
mediation. *Processes*, 11(4):1274, 2023.
- Clark, D. The design philosophy of the darpa internet pro-
tocols. In *Symposium proceedings on Communications*
architectures and protocols, pp. 106–114, 1988.
- Crispin, M. Internet message access protocol - version
4rev1. RFC 3501, March 2003. <https://www.rfc-editor.org/rfc/rfc3501>.
- Drouin, A., Gasse, M., Caccia, M., Laradji, I. H., Del Verme,
M., Marty, T., Boisvert, L., Thakkar, M., Cappart, Q.,
Vazquez, D., et al. Workarena: How capable are web
agents at solving common knowledge work tasks? *arXiv*
preprint arXiv:2403.07718, 2024.
- Eclipse Foundation. Eclipse Imos: Language model operat-
ing system. <https://eclipse.dev/lmos/>, 2025.
Accessed: 2025-05-20.
- Ehtesham, A., Singh, A., Gupta, G. K., and Kumar, S. A
survey of agent interoperability protocols: Model con-
text protocol (mcp), agent communication protocol (acp),
agent-to-agent protocol (a2a), and agent network protocol
(anp). *arXiv preprint arXiv:2505.02279*, 2025.
- Eiras, F., Petrov, A., Vidgen, B., De Witt, C. S., Pizzati,
F., Elkins, K., Mukhopadhyay, S., Bibi, A., Csaba, B.,
Steibel, F., et al. Near to mid-term risks and oppor-
tunities of open-source generative ai. *arXiv preprint*
arXiv:2404.17047, 2024.
- Elkhodr, M., Shahrestani, S., and Cheung, H. The internet
of things: New interoperability, management and security
challenges. *arXiv preprint arXiv:1604.04824*, 2016.
- Ferber, J. and Weiss, G. *Multi-agent systems: an intro-*
duction to distributed artificial intelligence, volume 1.
Addison-wesley Reading, 1999.
- Feuerriegel, S., Hartmann, J., Janiesch, C., and Zschech, P.
Generative ai. *Business & Information Systems Engineer-*
ing, 66(1):111–126, 2024.
- Fielding, R. T. and Reschke, J. Hypertext transfer protocol
(http/1.1): Message syntax and routing. <https://www.rfc-editor.org/rfc/rfc7230>, 2014. RFC
7230.
- Finin, T., Fritzson, R., McKay, D., and McEntire, R. Kqml
as an agent communication language. In *Proceedings*
of the third international conference on Information and
knowledge management, pp. 456–463, 1994.

- 495 FIPA. Fipa 97 specification, part 2: Agent communica-
496 tion language. Technical report, Foundation for Intelli-
497 gent Physical Agents, Geneva, Switzerland, 1997. URL
498 <http://www.fipa.org/specs/fipa00018/>.
499
- 500 FIPA. Fipa acl message structure specification. Technical
501 Report SC00061G, Foundation for Intelligent Physical
502 Agents, Geneva, Switzerland, 2002. URL <http://www.fipa.org/specs/fipa00061/>.
503
- 504 Ghareeb, A. E., Chang, B., Mitchener, L., Yiu, A.,
505 Szostkiewicz, C. J., Laurent, J. M., Razzak, M. T., White,
506 A. D., Hinks, M. M., and Rodrigues, S. G. Robin: A multi-
507 agent system for automating scientific discovery, 2025.
508 URL <https://arxiv.org/abs/2505.13400>.
509
- 510 Ghosh, D. P. Agentic ecosystem in engineering design: A
511 framework for interoperable legacy tools and emergent
512 collaboration via mcp/a2a protocols. 2025.
513
- 514 Google DeepMind. Announcing the
515 agent2agent protocol (a2a). [https://](https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/)
516 [developers.googleblog.com/en/](https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/)
517 [a2a-a-new-era-of-agent-interoperability/](https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/),
518 April 2025. Accessed: 2025-05-20.
519
- 520 Gürcan, Ö. Llm-augmented agent-based modelling for so-
521 cial simulations: Challenges and opportunities. *HAI*
522 *2024: Hybrid Human AI Systems for the Social Good*, pp.
523 134–144, 2024.
524
- 525 Hardt, D. The oauth 2.0 authorization framework. Technical
526 report, 2012.
527
- 528 Hosseini, S. and Seilani, H. The role of agentic ai in shaping
529 a smart future: A systematic review. *Array*, pp. 100399,
530 2025.
- 531 Hughes, L., Dwivedi, Y. K., Malik, T., Shawosh, M., Al-
532 bashrawi, M. A., Jeon, I., Dutot, V., Appanderanda, M.,
533 Crick, T., De’, R., et al. Ai agents and agentic systems: A
534 multi-expert analysis. *Journal of Computer Information*
535 *Systems*, pp. 1–29, 2025.
536
- 537 Kapoor, S., Kolt, N., and Lazar, S. Resist platform-
538 controlled ai agents and champion user-centric agent ad-
539 vocates. *arXiv preprint arXiv:2505.04345*, 2025.
- 540 Klensin, J. Simple mail transfer protocol. RFC 5321, Octo-
541 ber 2008. [https://www.rfc-editor.org/rfc/](https://www.rfc-editor.org/rfc/rfc5321)
542 [rfc/](https://www.rfc-editor.org/rfc/rfc5321)
543 [rfc5321](https://www.rfc-editor.org/rfc/rfc5321).
544
- 545 Kolos, S. Corba: a practical introduction. 2000.
546
- 547 Lampson, B. W. Hints for computer system design. In
548 *Proceedings of the ninth ACM symposium on Operating*
549 *systems principles*, pp. 33–48, 1983.
- Leiner, B. M., Cerf, V. G., Clark, D. D., Kahn, R. E., Klein-
rock, L., Lynch, D. C., Postel, J., Roberts, L. G., and
Wolff, S. S. The past and future history of the internet.
Communications of the ACM, 40(2):102–108, 1997.
- Leiner, B. M., Cerf, V. G., Clark, D. D., Kahn, R. E., Klein-
rock, L., Lynch, D. C., Postel, J., Roberts, L. G., and
Wolff, S. A brief history of the internet. *ACM SIGCOMM*
computer communication review, 39(5):22–31, 2009.
- Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V.,
Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel,
T., et al. Retrieval-augmented generation for knowledge-
intensive nlp tasks. *Advances in neural information pro-
cessing systems*, 33:9459–9474, 2020.
- Li, N., Gao, C., Li, M., Li, Y., and Liao, Q. Econa-
gent: large language model-empowered agents for
simulating macroeconomic activities. *arXiv preprint*
arXiv:2310.10436, 2023.
- Liu, B., Li, X., Zhang, J., Wang, J., He, T., Hong, S., Liu,
H., Zhang, S., Song, K., Zhu, K., et al. Advances and
challenges in foundation agents: From brain-inspired in-
telligence to evolutionary, collaborative, and safe systems.
arXiv preprint arXiv:2504.01990, 2025.
- Liu, Y., Jia, Y., Geng, R., Jia, J., and Gong, N. Z. For-
malizing and benchmarking prompt injection attacks and
defenses. In *33rd USENIX Security Symposium (USENIX*
Security 24), pp. 1831–1847, 2024.
- Lu, C., Lu, C., Lange, R. T., Foerster, J., Clune, J., and Ha,
D. The ai scientist: Towards fully automated open-ended
scientific discovery. *arXiv preprint arXiv:2408.06292*,
2024.
- Malka, M., Farhan, E., Morgenstern, H., and Shlezinger,
N. Decentralized low-latency collaborative inference via
ensembles on the edge. *IEEE Transactions on Wireless*
Communications, 2024.
- Marro, S., La Malfa, E., Wright, J., Li, G., Shadbolt, N.,
Wooldridge, M., and Torr, P. A scalable communication
protocol for networks of large language models. *arXiv*
preprint arXiv:2410.11905, 2024.
- Microsoft. Asynchronous request-reply pattern, 2025.
URL [https://learn.microsoft.com/](https://learn.microsoft.com/en-us/azure/architecture/patterns/async-request-reply)
[en-us/azure/architecture/patterns/](https://learn.microsoft.com/en-us/azure/architecture/patterns/async-request-reply)
[async-request-reply](https://learn.microsoft.com/en-us/azure/architecture/patterns/async-request-reply). Accessed: 2025-05-21.
- Model Context Protocol Project. Model context protocol
specification. [https://modelcontextprotocol.](https://modelcontextprotocol.io/specification/2025-03-26)
[io/specification/2025-03-26](https://modelcontextprotocol.io/specification/2025-03-26), 2025. Ac-
cessed: 2025-05-20.

- 550 Mogul, J. C. Clarifying the fundamentals of http. In *Proceedings of the 11th international conference on World*
551 *Wide Web*, pp. 25–36, 2002.
- 552
553
- 554 Murley, P., Ma, Z., Mason, J., Bailey, M., and Kharraz, A.
555 Websocket adoption and the landscape of the real-time
556 web. In *Proceedings of the Web Conference 2021*, pp.
557 1192–1203, 2021.
- 558
- 559 Murugesan, S. The rise of agentic ai: Implications, concerns,
560 and the path forward. *IEEE Intelligent Systems*, 40(2):
561 8–14, 2025.
- 562
- 563 Naptha AI. Naptha ai. <https://naptha.ai/>, 2025.
564 Accessed: 2025-05-22.
- 565
- 566 NEAR AI. Agent interaction and transaction protocol (aitp)
567 specification. <https://aitp.dev/>, 2025. Accessed:
568 2025-05-20.
- 569
- 570 OpenAI. Deep research system card. [https://cdn.](https://cdn.openai.com/deep-research-system-card.pdf)
571 [openai.com/deep-research-system-card.](https://cdn.openai.com/deep-research-system-card.pdf)
572 [pdf](https://cdn.openai.com/deep-research-system-card.pdf), 2025.
- 573
- 574 Paniagua, C. Service interface translation. an interoperabil-
575 ity approach. *Applied Sciences*, 11(24):11643, 2021.
- 576
- 577 Papazoglou, M. P. and Van Den Heuvel, W.-J. Service
578 oriented architectures: approaches, technologies and re-
579 search issues. *The VLDB journal*, 16:389–415, 2007.
- 580
- 581 Park, J. S., O’Brien, J., Cai, C. J., Morris, M. R., Liang,
582 P., and Bernstein, M. S. Generative agents: Interactive
583 simulacra of human behavior. In *Proceedings of the 36th*
584 *Annual ACM Symposium on User Interface Software and*
585 *Technology*, UIST ’23, New York, NY, USA, 2023. Asso-
586 ciation for Computing Machinery. ISBN 9798400701320.
587 doi: 10.1145/3586183.3606763. URL [https://doi.](https://doi.org/10.1145/3586183.3606763)
588 [org/10.1145/3586183.3606763](https://doi.org/10.1145/3586183.3606763).
- 589
- 590 Peng, W. and Cisna, J. Http cookies—a promising technology.
591 *Online Information Review*, 24(2):150–153, 2000.
- 592
- 593 Perplexity. Perplexity.ai. <https://perplexity.ai/>,
594 2025. Accessed: 2025-05-22.
- 595
- 596 Postel, J. Transmission control protocol. RFC 793, Septem-
597 ber 1981. [https://www.rfc-editor.org/rfc/](https://www.rfc-editor.org/rfc/rfc793)
598 [rfc793](https://www.rfc-editor.org/rfc/rfc793).
- 599
- 600 Raymond, E. The cathedral and the bazaar. *Knowledge,*
601 *Technology & Policy*, 12(3):23–49, 1999.
- 602
- 603 Reed, D., Sporny, M., Longley, D., Allen, C., Grant, R.,
604 Sabadello, M., and Holt, J. Decentralized identifiers
605 (dids) v1. 0. *Draft Community Group Report*, 2020.
- 606
- 607 Research, A. The danger of fragmenta-
608 tion in ai agent ecosystems, 2024. URL
609 [https://www.arionresearch.com/blog/](https://www.arionresearch.com/blog/x8nwy1u8hufvp84sczuskkg49qspop)
610 [x8nwy1u8hufvp84sczuskkg49qspop](https://www.arionresearch.com/blog/x8nwy1u8hufvp84sczuskkg49qspop). Accessed:
611 2025-05-21.
- 612
- 613 Ruan, J., Chen, Y., Zhang, B., Xu, Z., Bao, T., Mao, H., Li,
614 Z., Zeng, X., Zhao, R., et al. Tptu: Task planning and
615 tool usage of large language model-based ai agents. In
616 *NeurIPS 2023 Foundation Models for Decision Making*
617 *Workshop*, 2023.
- 618
- 619 Russell, S. J. and Norvig, P. *Artificial intelligence: a modern*
620 *approach*. pearson, 2016.
- 621
- 622 Sapkota, R., Roumeliotis, K. I., and Karkee, M. Ai agents
623 vs. agentic ai: A conceptual taxonomy, applications and
624 challenge. *arXiv preprint arXiv:2505.10468*, 2025.
- 625
- 626 Schneider, J. Generative to agentic ai: Survey, conceptual-
627 ization, and challenges. *arXiv preprint arXiv:2504.18875*,
628 2025.
- 629
- 630 Schwartz, C. Web search engines. *Journal of the American*
631 *Society for Information Science*, 49(11):973–982, 1998.
- 632
- 633 Shafer, L. R. Open standards and the digital age: History,
634 ideology, and networks. *Afterimage*, 42(3):39, 2014.
- 635
- 636 Shavit, Y., Agarwal, S., Brundage, M., Adler, S., O’Keefe,
637 C., Campbell, R., Lee, T., Mishkin, P., Eloundou, T.,
638 Hickey, A., et al. Practices for governing agentic ai sys-
639 tems. *Research Paper, OpenAI*, 2023.
- 640
- 641 Shlezinger, N., Farhan, E., Morgenstern, H., and Eldar, Y. C.
642 Collaborative inference via ensembles on the edge. In
643 *ICASSP 2021-2021 IEEE International Conference on*
644 *Acoustics, Speech and Signal Processing (ICASSP)*, pp.
645 8478–8482. IEEE, 2021.
- 646
- 647 The Linux Foundation. Linux Foundation Launches
648 the Agent2Agent Protocol Project to Enable Secure,
649 Intelligent Communication Between AI Agents. [https:](https://www.linuxfoundation.org/press/linux-foundation-launches-the-agent2agent-protocol)
650 [www.linuxfoundation.org/press/](https://www.linuxfoundation.org/press/linux-foundation-launches-the-agent2agent-protocol)
651 [linux-foundation-launches-the-agent2agent-protoco](https://www.linuxfoundation.org/press/linux-foundation-launches-the-agent2agent-protocol)
652 [l](https://www.linuxfoundation.org/press/linux-foundation-launches-the-agent2agent-protocol)
653 [2025a](https://www.linuxfoundation.org/press/linux-foundation-launches-the-agent2agent-protocol). Accessed: 2026-01-29.
- 654
- 655 The Linux Foundation. ACP Joins Forces
656 with A2A. [https://lfaidata.](https://lfaidata.foundation/communityblog/2025/08/29/acp-joins-forces-with-a2a-under-the-linux-foundation)
657 [foundation/communityblog/2025/08/29/](https://lfaidata.foundation/communityblog/2025/08/29/acp-joins-forces-with-a2a-under-the-linux-foundation)
658 [acp-joins-forces-with-a2a-under-the-linux-foundat](https://lfaidata.foundation/communityblog/2025/08/29/acp-joins-forces-with-a2a-under-the-linux-foundation)
659 [ion](https://lfaidata.foundation/communityblog/2025/08/29/acp-joins-forces-with-a2a-under-the-linux-foundation)
660 [2025b](https://lfaidata.foundation/communityblog/2025/08/29/acp-joins-forces-with-a2a-under-the-linux-foundation). Accessed: 2026-01-29.
- 661
- 662 Tran, K.-T., Dao, D., Nguyen, M.-D., Pham, Q.-V.,
663 O’Sullivan, B., and Nguyen, H. D. Multi-agent collab-
664 oration mechanisms: A survey of llms. *arXiv preprint*
665 *arXiv:2501.06322*, 2025.

- 605 W3C AI Agent Protocol Community Group. AI Agent
606 Protocol Specification. <https://w3c-cg.github.io/ai-agent-protocol/>, 2025. Accessed: 2026-
607 01-29.
- 608
609
- 610 Wei, A., Haghtalab, N., and Steinhardt, J. Jailbroken: How
611 does llm safety training fail? *Advances in Neural Informa-*
612 *tion Processing Systems*, 36:80079–80110, 2023.
- 613
614 Wildcard AI. agents.json specification. <https://docs.wild-card.ai/agentsjson/introduction>,
615 2025. Accessed: 2025-05-20.
- 616
- 617 Yan, X. and Feng, Y. Mandating interoperability of in-
618 terpersonal communications services in an ecosystems
619 competition context. *Computer Law & Security Review*,
620 50:105850, 2023.
- 621
- 622 Zamyatin, A., Al-Bassam, M., Zindros, D., Kokoris-Kogias,
623 E., Moreno-Sanchez, P., Kiayias, A., and Knottenbelt,
624 W. J. Sok: Communication across distributed ledgers.
625 In *Financial Cryptography and Data Security: 25th In-*
626 *ternational Conference, FC 2021, Virtual Event, March*
627 *1–5, 2021, Revised Selected Papers, Part II 25*, pp. 3–36.
628 Springer, 2021.
- 629
630 Zhao, Q., Wang, Y., Yang, B., Shang, K., Sun, M., Wang,
631 H., Yang, Z., and Xin, H. A comprehensive overview of
632 security vulnerability penetration methods in blockchain
633 cross-chain bridges. *Authorea Preprints*, 2023.
- 634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659

A. Ongoing Efforts

Our work complements several concurrent efforts toward agentic AI standardization. The W3C AI Agent Protocol Community Working Group is developing specifications for agent communication and interoperability within the web standards ecosystem (W3C AI Agent Protocol Community Group, 2025). The AGNTCY project (Agency.org, 2025) provides vendor-agnostic infrastructure for agent discovery, identity, messaging, and observability, enabling cross-framework collaboration. Notably, The Linux Foundation has taken a leadership role in fostering interoperability by hosting both the A2A and ACP protocols under its governance structure, signaling industry commitment to open, collaborative standards development (The Linux Foundation, 2025a;b). Additionally, industry consortia and academic initiatives are exploring security frameworks, trust models, and ethical guidelines for multi-agent systems. Rather than competing with these efforts, our minimal standards approach provides a technical foundation that can accommodate diverse governance models and specialized extensions. We view WEB OF AGENTS as a baseline architecture upon which these complementary initiatives can build, ensuring that efforts toward security, discovery infrastructure, and regulatory compliance operate within an interoperable ecosystem rather than perpetuating fragmentation.

B. Implementation Details

To demonstrate the feasibility of our proposed WEB OF AGENTS blueprint, we implemented the four building blocks (agent-to-agent messaging, interaction interoperability, state management, and agent discovery) as extensions to existing agentic AI frameworks. Specifically, we forked and modified the codebases of A2A, ACP, and Agora to incorporate HTTP-based messaging, interaction documentation endpoints, session-based state management, and capability advertisement at well-known paths. Our open-source proof-of-concept implementation required minimal modifications to each framework—typically under 200 lines of code per protocol—demonstrating that adoption of these standards imposes negligible engineering overhead². Furthermore, we ensure that agents within an ecosystem, for example, two agents in the A2A ecosystem, can still operate as intended by the framework developers. The demo showcases cross-ecosystem agent collaboration where agents implemented in different protocols can discover, interact with, and maintain stateful conversations with one another.

²Source code available at https://anonymous.4open.science/r/web_of_agents_demo-B011/

C. Politics and Economics

The case for interoperability in agentic AI extends beyond technical architecture into political and economic domains. Questions of ecosystem fragmentation, vendor lock-in, market concentration, and user autonomy are fundamentally about power distribution, competitive dynamics, and who controls the infrastructure of collaborative AI systems. We acknowledge that interoperability is not merely a technical concern but a sociotechnical one with significant implications for innovation, competition, and access. However, in this position paper, we deliberately take a technical stand. We focus on demonstrating that minimal interoperability standards are architecturally feasible, practically implementable, and sufficient for enabling cross-ecosystem collaboration. This technical foundation is necessary for addressing the broader political and economic considerations. By establishing that interoperability can be achieved through lightweight extensions to existing web technologies, we provide concrete evidence that the technical barriers are surmountable. The political and economic arguments for or against interoperability, while crucial, must be informed by an understanding of what technical solutions are viable and what tradeoffs they entail.