

Intelligent Agents: Where Is the Web?

Victor Charpenay

LIMOS, Mines Saint-Etienne

October 26th, 2025

Agents and the Semantic Web

The Semantic Web


Agents and the Semantic Web

James Hendler, *University of Maryland*

At a colloquium I attended recently, a speaker described a “science fiction” vision comprising agents running around the Web performing complex actions for their users. The speaker argued that we are far from the day this vision would become a reality because we don’t have the infrastructure to make it happen.

Although I agree with his assessment about infrastructure, his claim that we are “far from the day” is too pessimistic. A crucial component of this infrastructure, a standardized Web ontology language, is emerging. This article offers a few pointers to this emerging area and shows how the ontology languages of the Semantic Web can lead directly to more powerful agent-based approaches—that is, to the realization of my colleague’s “science fiction” vision.

Many challenges of bringing communicating multiagent systems to the Web require ontologies.



J. Hendler (2001). Agents and the Semantic Web, IEEE Intelligent Systems, IEEE.

Agents and the Semantic Web?



Editor: James Hendler
Rensselaer Polytechnic Institute
hendler@cs.rpi.edu

A Letter from the Editor

Where Are All the Intelligent Agents?

Intelligent Readers,

In the late 1990s, many of us believed that the large-scale deployment of “intelligent agent”-based computing was right around the corner. In the early 2000s, much US and international research funding focused on making this happen. This magazine and many others had hugely popular special issues on agents, and academic conferences on agent-based computing were abuzz, with many of us “old timers” beginning to believe that the time had really come. But now, looking at what’s hot on the Web, in IT development, and in venture capital circles, I find myself shaking my head and wondering, “Where are all the agents?”

J. Hendler (2007). Where Are All the Intelligent Agents?, IEEE Intelligent Systems, IEEE.

Here Are Intelligent Agents

Compile a research report on how the retail industry has changed in the past 3 years.

Searched for Statista e-commerce share

I reviewed e-commerce share data for the US from 2019 to 2023 noting a peak around 2020 and stabilization in the subsequent years. Also, I looked into global e-commerce trends, particularly China's significant influence.

Read oberlo.com

OpenAI consistently upholds a dedication to ethical research and the well-being of all individuals.

Searched for global e-commerce share 2019 vs. 2020

Checking for trends in e-commerce share growth from 2019 to 2020.

OpenAI (2025). Introducing Deep Research, openai.com.

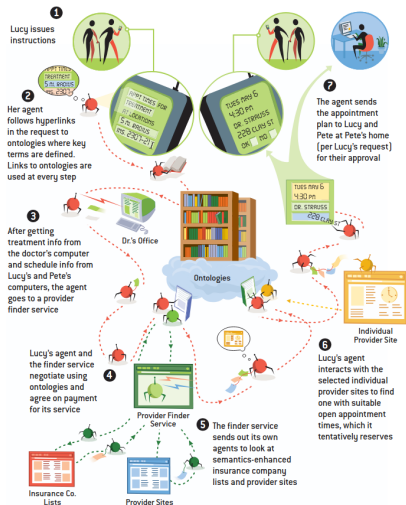
What is a Web Task?

Definition

A task performed by an intelligent agent is specific to the Web if all of the following criteria hold:

1. the client reads the representation of resources from at least two servers during execution;
2. the client reads the representation a resource only if a link to that resource exists in the representation of a previously accessed resource or it is the first resource being read;
3. the client writes the representation of a resource only if a form describing the operation exists in the representation of a previously accessed resource;
4. the client executes no server-sent code (optional).

Use Case (Service Composition)



- ▶ Each participant in the task has its own server, for instance a Social Linked Data (Solid) Pod
- ▶ Each participant has its own “steward” agent that has privileged access to their data but not others’

T. Berners-Lee, J. Hendl, O. Lassila (2001). The Semantic Web, Scientific American, Springer Nature.

Use Case (Service Selection)

Get me a satellite photo of this region of the Atlantic.

[map with bounding box]

- ▶ A satellite image taken yesterday at 10 AM is available on the Web at [http://...](#)
- ▶ A new satellite image, to be taken today at 10 AM, will be available for \$100—click here to authorize transfer of funds and obtain image. (You will need a valid credit card number from one of the following providers ...)
- ▶ In an emergency situation, a Coast Guard observer plane can be sent to any location within the area you indicate. Service Note: You will be responsible for cost of flight if the situation does not result in an emergency pickup. Click here for more information.
- ▶ A high-altitude observer can be sent to your location in 13 hours. Click here to initiate procedure. (You will need to provide US military authorization, a valid military unit code, and the name of the commanding officer. Abuse of this procedure can result in fine or imprisonment.)
- ▶ A service entitled commercial service for providing satellite images is advertised as becoming available in 2004. See [http://...](#) for more information.

Agent Architectures

Table: Main agent architecture in the literature

Architecture	Pub. date	Description
ReAct	2023	Reasoning and acting
Toolformer	2023	Tool usage
MindAct	2023	Ranking before acting
SeeAct	2024	Perception via text and images
SeePlanAct	2024	Planning and memory

System Prompt

You are an expert assistant who can solve any task using tool calls. You will be given a task to solve as best you can.

To do so, you have been given access to some tools.

The tool call you write is an action: after the tool is executed, you will get the result of the tool call as an "observation". This Action/Observation can repeat N times, you should take several steps when needed.

You can use the result of the previous action as input for the next action.

You have access to these tools:

code-generated content

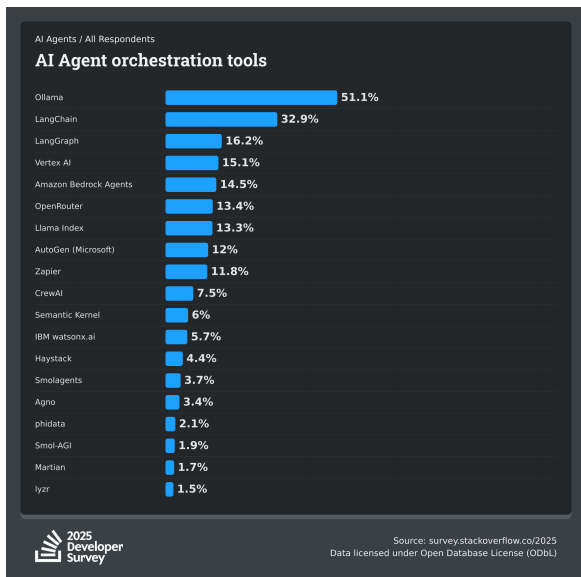
Now Begin!

Agent Frameworks

Table: Popular agent frameworks based on the ReAct design pattern (column 'Count' gives the number of tools related to the Web over the total number of tools built into the framework)

Framework	Maintainer	Count	Example Web tools
LangGraph	LangChain	26/35	Google search, Tavily search, Slack, Hyberbrowser, ...
CrewAI	CrewAI	21/52	Google Serper, Firecrawl, Selenium, Website RAG search, ...
AG2	Microsoft, UPenn	9/15	Google search, Perplexity search, Firecrawl, deep research, ...
Smolagents	HuggingFace	5/9	Google search, API Web search, visit Web page, ...
Agents SDK	OpenAI	1/7	Web search

More Agent Frameworks



Benchmark Datasets

Question answering: the agent must find evidence anywhere on the Web to answer a question.

Browser Automation: the agent must perform a task on one or more identified websites.

Question Answering Datasets

Table: Most frequent Web sites used as sources in various question answering datasets: SimpleQA (evaluation split), GPQA (diamond split) and GAIA (not all tasks reviewers explicitly mentioned their sources, figures are likely underestimated)

ELI5 (KILT version)	SimpleQA	GPQA	GAIA
wikipedia.org (100%)	wikipedia.org (40%) fandom.com (3%) ...	wikipedia.org (18%) nih.gov (10%) libretexts.org (8%) stackexchange.com (3%) ...	wikipedia.org (18%) youtube.com (5%) github.com (3%) researchgate.net (3%) ...

ELI5 and SimpleQA only require a single source.

Browser Automation Datasets

Table: Benchmark datasets used to evaluate agents in a Web environment (column 'X' gives the ratio of cross-site navigation tasks)

Dataset	X	Percepts	Actions
QAWoB	0%	RGB image, DOM	select date, drop down, scroll, click
FormWoB	0%	RGB image, DOM	click, drag, scroll, copy, paste, select
MiniWoB++	0%	RGB image, DOM	move, click, drag, scroll, copy, paste, select
WebShop	0%	page type, text	search, choose
Mind2Web	0%	RGB image (<i>optional</i>), HTML	click, select
WebArena	5.9%	RGB image, HTML, text	click, scroll, go back/forward, go to
WebLIX	20%	RGB image, HTML	click, load, say, submit
AssistantBench	61%	RGB image, HTML	scroll, go to, search, go back

The best agent achieves 44% of the tasks on WebArena but only 2% of the tasks on AssistantBench.

See the BrowserGym Leaderboard for latest evaluations.

Alternative Proposal

An agent needs to find out what art museums are located in Pittsburgh by searching Wikipedia. Next, it should identify the location of each museum on a map, optimizing the itinerary based on the information collected. Finally, the agent needs to update the README file in the appropriate repository with the planned route.

```
SELECT ?museum ?long ?lat
WHERE {
    ?museum dbo:location dbr:Pittsburgh ; a dbo:Museum ;
        geo:long ?long ; geo:lat ?lat .
}
```

Embedded RDF in HTML




Figure: Statistics of RDF triples about non-Wikipedia resources

More than 80% of all resources involved in ELI5, SimpleQA, GPQA and GAIA have an RDF description.

Conclusion

Many content providers differentiate software agents from humans on the Web.

Many publish machine-readable, easier-to-process versions of their content.

A Web agent should answer queries a single source cannot answer; achieve tasks a single platform cannot achieve.