

25 years ago...



EXCLUSIVE: WARP DRIVE UNDERWATER • ARCTIC OIL VS. WILDLIFE

SCIENTIFIC AMERICAN

MAY 2001 \$4.95
WWW.SCIAM.COM

Get the Idea? (TOMORROW'S WEB WILL)




PLUS:
Antibiotics'
Dim Future
Rorschach:
A Waste of Ink
The Oldest Stars

25 years ago...

We published an article:

Berners-Lee et al: “The Semantic
Web”, Scientific American
284(5):34–43 (May 2001)



Is the Semantic Web vision still relevant?

Dr. James Hendler

Professor

Rensselaer Polytechnic Institute

Dr. Ora Lassila

Assoc. Dir. of Data Engineering & Governance

Accenture

Knowledge Graph Conference, May 2026

Yes!


Just one minor problem...

We knew **what we had**

all the symbolic AI stuff: KR, agents, planning, etc.

We knew **what we wanted**

But how to get there...?



Mixed reception

Some people got very excited; over the years the article has **motivated a lot of amazing work**

Some people thought we took drugs

We were called us “sci-fi authors”

Many claims over the years that the Semantic Web is generally a failure



We think the joke is now on the critics

The Semantic Web & associated standards (RDF, OWL, SHACL, SPARQL) have matured into modern enterprise knowledge graphs

Majority of all Web pages contain some embedded RDF

All Adobe documents contain RDF as metadata

*“RDF has turned out to be the most
catastrophically successful failure there is.”*

– Charles Ivie (Ortecha, ex-AWS)

Timeline

The miracle happened!

mid-1990s	Tim's original writings, term "Semantic Web" coined
1997-1999	First standards for machine-interpretable data on the Web (RDF)
1999	US White House calls for research in "intelligent agents"
2000	DARPA Agent Markup Language program ("DAML"), DAML+OIL
2001	Scientific American article, ISWC conference series started
2004	OWL
2008	SPARQL, ISWC Billion Triple Challenge
2011	schema.org
2012	Google "Knowledge Graph" publicized
2017	Semantic markup on majority of Web content
late 2010s	KGs move to mainstream, enterprise-ready Semantic Web tools emerge

Mainstream?

Since 2018 (or so), enterprise knowledge graphs have become “acceptable” and increasingly popular

More serious (“enterprise-ready”) Semantic Web tools

KGC started in 2019

What has happened with knowledge graphs and ontologies follows a typical pattern observed with AI technologies for a long time

Typical pattern?

Generative AI and LLMs have transformed how we apply AI

However...

Possibility of a 3rd “AI Winter” looms

we must manage expectations (failure to do so led to the previous winter)

Old criticisms of AI are suddenly very relevant today

ELIZA (Weizenbaum 1966)

Chinese Room (Searle 1980)

Agents: Back to the future

The original Semantic Web vision was really about **agents**

Transition towards software that **works on users' behalf**, not as tools

The SciAm article was published during the tail end of the “AI winter”, thus the parts of the story that “smacked of AI” were promptly forgotten

With the advent of generative AI and LLMs, the idea of an “agent” has been revived

Agents: Back to the future

The classical idea of an autonomous agent requires several symbolic AI techniques to implement

The advent of the WWW was much welcomed by those of us who aspired to build agents in the 1990s and earlier

Historical challenge: user interaction (natural language, etc.)

Agentic AI?

Is “agentic” even a word?

Contemporary “agentic AI” is not really what we imagined earlier

good news: we can now converse with our agents

Bad news: current solutions do not generalize

we are still just using pre-scripted workflow, process automation, triggered execution of scripts, etc. – and for a limited purpose

The next miracle...?

Just one minor problem...


We know what we **have**

enterprise-grade KGs & ontologies
generative AI (LLMs)

We know what we **want**

**the original Semantic Web vision: real
“intelligent” agents on the Web**

But how to get there...?



Do we really understand what “agency” is?

agency (n): the ability to make decisions and act independently on someone else's behalf

Someone who has agency wields **your decision-making power**, or part thereof, to make decisions and to act – without consulting you

We are used to this when dealing with humans, but to build agents, we need to understand what the implications are

What do we need for “real” agents now?

Business

Access to information sources

Model of the world

Intentionality

Safe autonomy

Technology

Interoperability (RDF, etc.)

Web, KGs (KR, ontologies, etc.)

Planning, goal-oriented processing

Guardrails, SOLID, security & privacy

These are all classical AI techniques

What do we need for “real” agents **later**?

Business

Multi-agent coordination

Frictionless user interaction

Standards for AI technologies
on the Web

Technology

New foundations beyond game
theory

More UX research, a revival of
ubiquitous computing

Research into describing capabilities
and functionality (revival of DAML-S
for APIs?)

Why do we need real agency?

Agency done right will allow AI to do what AI does best

agents can do for AI what the Web did for hypertext

(question assumptions, provide interoperability and scalable foundations)

Agency done right will allow humans to do **what humans do best**

“see the big picture”: interact and collaborate to solve **big** problems

free us from mundane routine to use our innate creativity

This is a moment for human liberation!

Questions?

hendler@cs.rpi.edu

ora.y.lassila@accenture.com

This is a moment for human liberation!