# Semantic Web and AI: Can we finally realize the vision?

Dr. Ora Lassila
Principal Technologist, Amazon Neptune
Co-chair, W3C RDF-star WG

Keynote Address
International Semantic Web Conference, November 2024

aws

# Who am I, and what have I done?

**Current:**

Principal Technologist, Amazon Neptune (AWS)
Co-chair, W3C RDF-star Working Group

**Past:**

State Street, Pegasystems, Nokia,
MIT, CMU, Helsinki University of Technology

**Education:**

Ph.D (D.Sc) CS & AI, Helsinki University of Technology

Semantic Web vision

RDF

KR for NASA Deep Space 1

2 daughters

43k+ citations

Grand Prize of Usenix
Obfuscated C Code Contest

25 years ago I said this is about AI,
but nobody listened

# Game plan

1. Yesterday
2. Today
3. Tomorrow

# Episode 1

…where I talk about what has happened

Good judgement comes from experience.
Experience comes from bad judgement.
– unknown

# One fateful day in 1996 @ MIT

Per my recollection, this is what happened…

TimBL: *"So Ora, what do you think is wrong with the Web?"*

me: *"Well Tim, the Web is for humans, and that makes it hard to automate anything. I want autonomous agents."*

TimBL: *"That's it! How do we solve this?"*

me: *"Ugh… I don't know. Maybe we should see if knowledge representation could be used."*

*TimBL: "Please look into this."*

# The Semantic Web?

The original Semantic Web vision shifted emphasis from "human-interpretable" to **"machine-interpretable"** Web content

Knowledge Representation

+

Wait!
There's more!

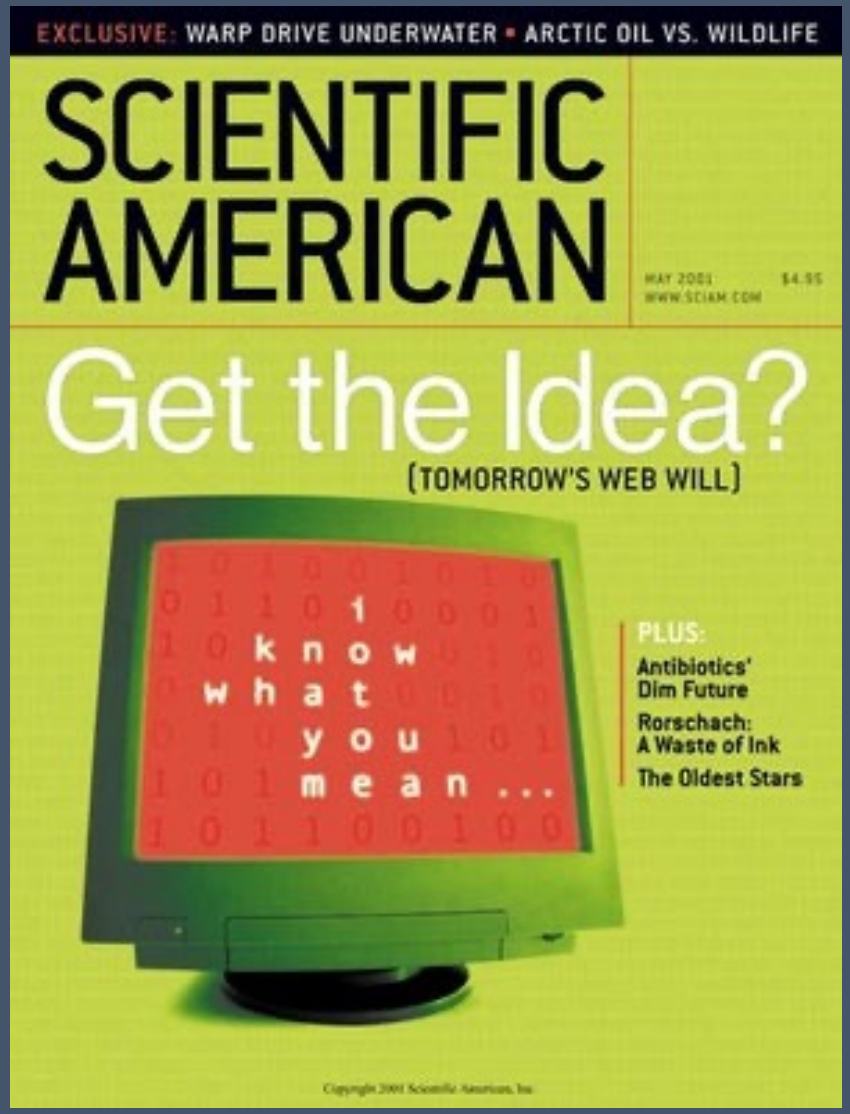

# The Semantic Web!

The vision is about humans interacting with software and devices, and enabling these to have **more autonomy**

Knowledge Representation + Autonomous Agents

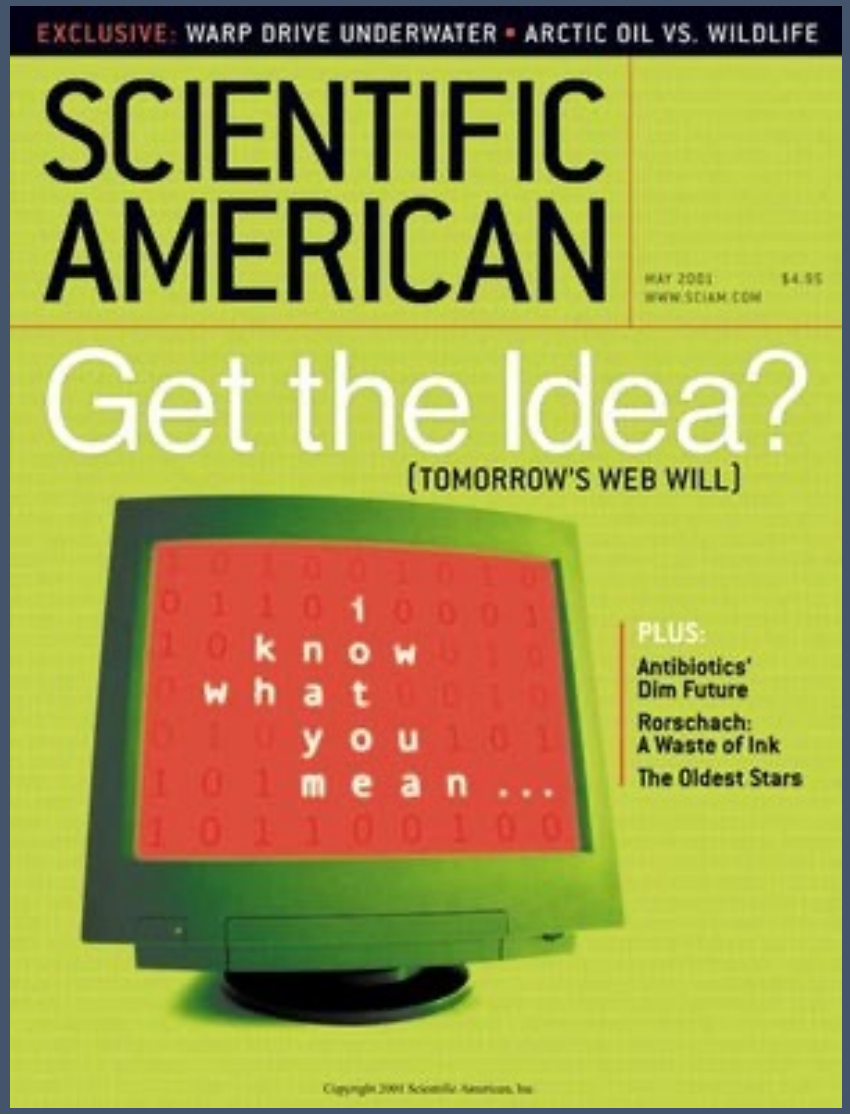

# Not just the Web, but also devices and the physical world

My own early use cases, experiments, prototypes, applications

- services, functionality, behavior (OWL-S)
- sensor data
- context-awareness
- policies

ubiquitous computing

- media (photos, music)
- communication with other users
- calendars & scheduling
- (and integration of all of the above)

mobile devices

# So far so good…

I found all this to be very exciting

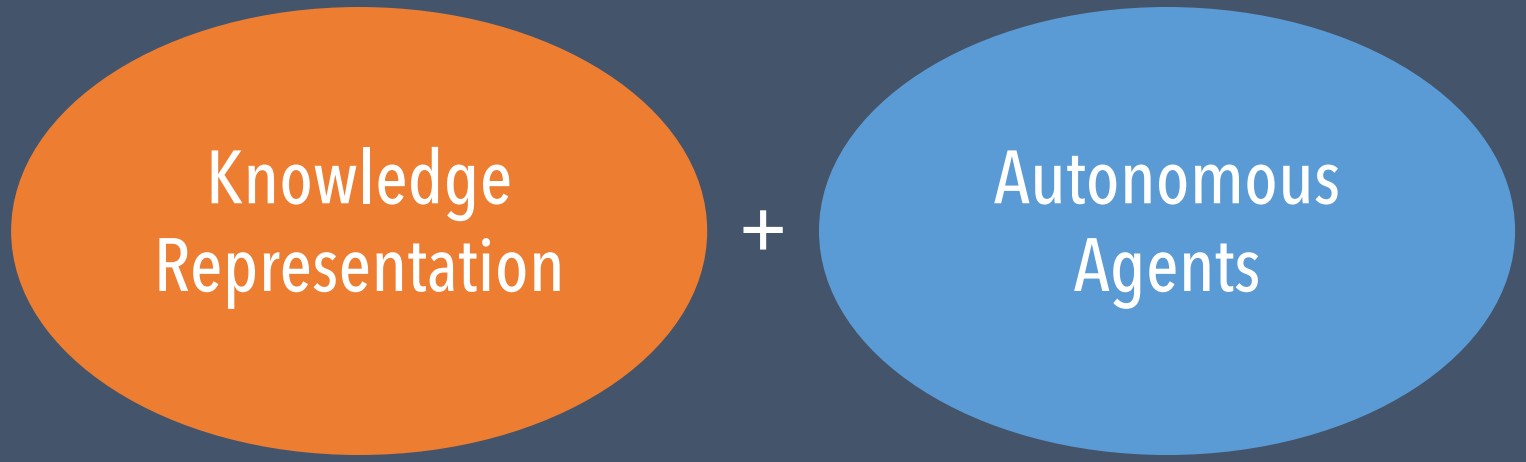


However…

# The vision was created in the middle of the "AI winter"…

…and people actively worked to remove AI from the Semantic Web

## Linked Data ≠ Semantic Web

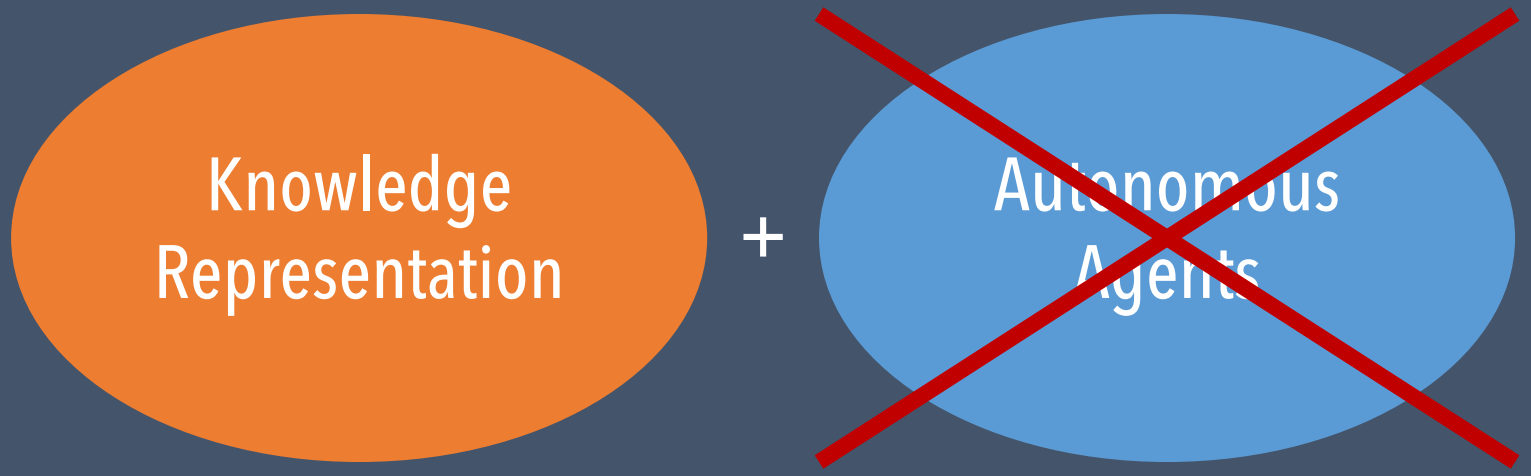


Kudos to Jim Hendler for naming DARPA's Semantic Web program the "DARPA Agent Markup Language"

# What about knowledge representation?

Reasoning is an important part of KR

Ontologies are typically built with some assumption of reasoning

Declarative vs. procedural?

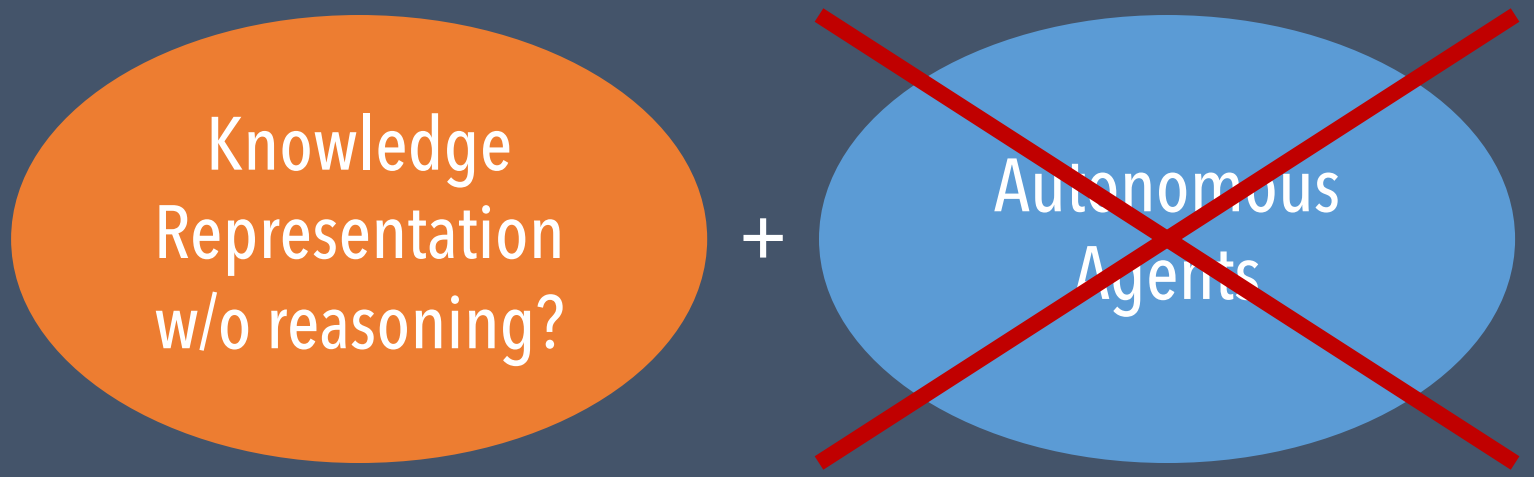

# …and, there were three things wrong with the term

# "The Semantic Web"

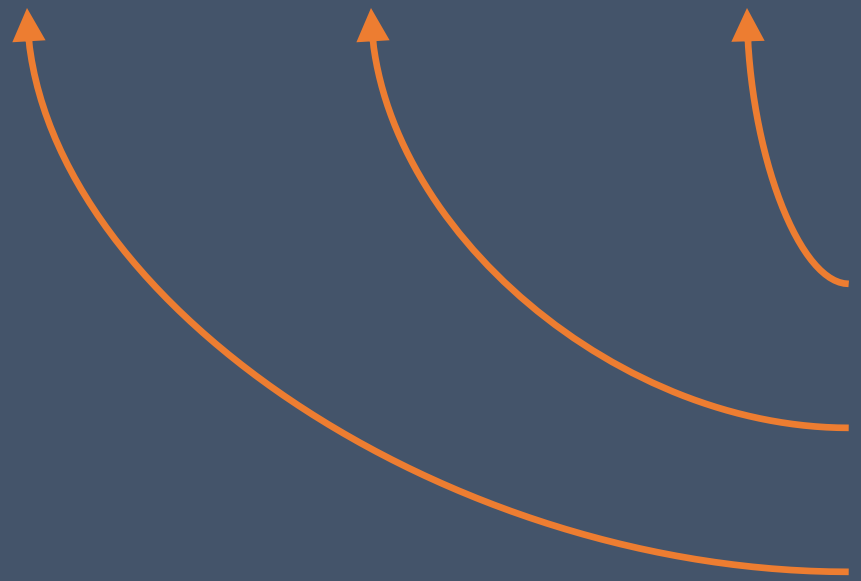

1. It is not (only) about the Web
2. Nobody knows what "semantic" means
3. There will not be just one

# Criticism and distractions…

“Why can’t I just do this with <some technology>?”

# Criticism and distractions…

"Why can't I just do this with <some technology>?"

Then: "Why can't I just do this with XML?" (this delayed us by years)

Recently: "Why can't I just do this with JSON?"

Today: "Why can't I just do this with Labeled Property Graphs?"

…

# Criticism and distractions…

"Why can't I just do this with <some technology>?"

"This is too academic…"

# Criticism and distractions…

"Why can't I just do this with <some technology>?"

"This is too academic…"

um… those people might actually be onto something

# Why do we show them this?

<https://somanyaircraft.com/data/colors/bs381c#175>
  <https://somanyaircraft.com/data/schemata/color/1/core#rgbHex>
  "4e83b6"^^<https://somanyaircraft.com/data/schemata/color/1/core#RGB> .
<https://somanyaircraft.com/data/colors/bs381c#175>
  <http://www.w3.org/1999/02/22-rdf-syntax-ns#type>
  <https://somanyaircraft.com/data/schemata/color/1/core#BS381CColor> .
<https://somanyaircraft.com/data/colors/bs381c#175>
  <http://www.w3.org/2000/01/rdf-schema#label>
  "BS381C 175 (Light French blue)" .
<https://somanyaircraft.com/data/colors/bs381c#175>
  <https://somanyaircraft.com/data/schemata/color/1/core#collection>
  <https://somanyaircraft.com/data/colors/bs381c#> .
<https://somanyaircraft.com/data/colors/bs381c#175>
  <https://somanyaircraft.com/data/schemata/color/1/core#bs381cCode>
  "175" .
<https://somanyaircraft.com/data/colors/bs381c#175>
  <https://somanyaircraft.com/data/schemata/color/1/core#name>
  "Light French blue" .
<https://somanyaircraft.com/data/colors/bs381c#175>
  <https://somanyaircraft.com/data/schemata/color/1/core#workingSpace>
  <https://somanyaircraft.com/data/schemata/color/1/core#sRGB> .

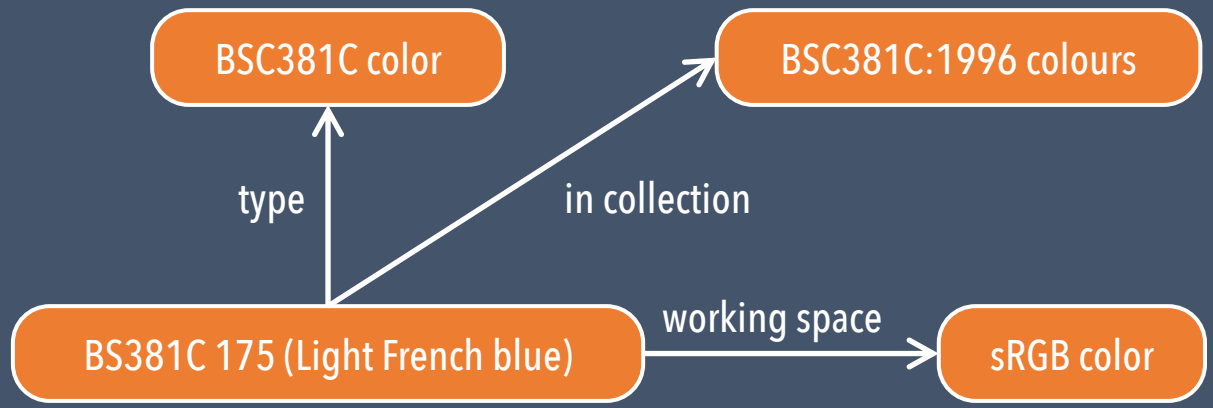


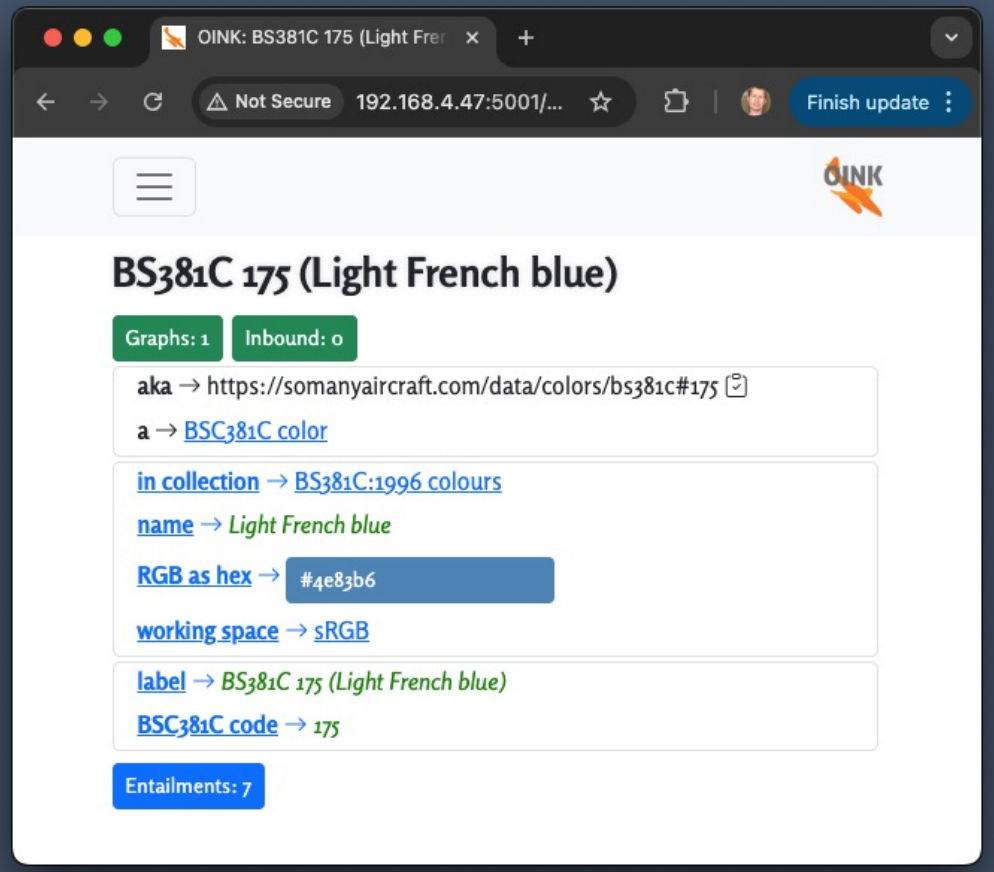


. . . when I can explain this to a 7-year old

# Criticism and distractions…

"Why can't I just do this with <some technology>?"

"This is too academic…"

"It does not scale"

# It does scale! Here is a great example…

**Inventory knowledge graph** of Amazon Fulfillment is used to

- investigate fulfillment processes
- investigate lost & found -issues
- improve precision of product recalls

Extends the PROV-O ontology: models the end-to-end logistics process as a form of provenance

Runs on multiple (federated) Neptune clusters

Size:

1T+ triples

4B new triples per day

Queries (p95):

< 50 ms to find a node

< 1 s to retrieve the whole path

# In summary…

An exciting, transformative idea!

…and yet, many people thought we were crazy, unrealistic, etc.

- *"the Scientific American article is science fiction"* (an actual quote)

(I think we all have proven those people wrong)

# Question I get sometimes:
# "Do you yourself use RDF?"

I do… for example, I organize my photos as a knowledge graph:

1. harvest XMP from images
2. annotate using SKOS schemes
3. create other datasets
4. collect notes as RDF

# Episode 2

…where I talk about where we are now
(and that we won)

# Semantic Web today?

Since the late 1990s, we have made a lot of progress

The Semantic Web, as it stands today, is the result of a lot of work by many smart people

So where, exactly, are we today?

# The Semantic Web today

The original vision has matured to modern enterprise knowledge graphs

Gartner says knowledge graphs, ontologies, RDF are a “major trend”

(so don’t take my word for it)

# The Semantic Web today

Someone asked me this years ago: “When can we declare victory?”

My answer: “When someone builds a knowledge graph -based system and does not feel the need to tell everyone that they have built a knowledge graph -based system”

I think we are there now

# The Semantic Web today

Also, everyone wants a “semantic layer” for their data stack…

(although I think many people are struggling to define what that is)

# The Semantic Web today

The limited realization of the original vision turned out to be useful

Knowledge graphs are great for

- data integration
- establishing business vocabularies
- insulating business applications from the technical details of (legacy) data

Knowledge graphs even have the quality that I always thought was the coolest thing about the Semantic Web: serendipity!

## Serendipity and “future-proofing”

The Semantic Web can be a solution to those problems and situations that we are yet to define

(yep, seriously)

# What did we achieve?

Standards:

- RDF, SPARQL, SHACL, R2RML, OWL, …

Vocabularies and ontologies:

- DC, DCAT, PROV-O, FHIR, FIBO, XMP, …

Software:

- libraries, reasoners, triple stores, …

Services and data:

- Wikidata, OpenStreetmap, …

# What did we miss?

Agents (not yet there)

Adoption (could be better)

Perception (there is work to be done)

Catalogs of ontologies and data

- (say "LOD Cloud" and you have to go to the back of the room)

Easier-to-use ontology editors

Good user experience

# What is different now? What would I do differently?

## "How would I write the SciAm article if I was writing it today?"

Reduction in scope (Web → enterprise) was a good thing

I had no idea how fast the volume of data would scale up

AI is acceptable today; not so much back then

(Also: I should have fought much harder to kill the XML syntax of RDF)

# Much of the original Semantic Web work came from W3C

The work still continues…

- RDF-star WG – new versions of RDF and SPARQL
- Autonomous Agents on the Web CG
- Knowledge Graph Construction CG
- Solid CG – new take on decentralization, “the Web, take 3”
- SPARQL-dev CG
- RDF Rust Common Crates CG
- Web of Things CG
- …

# Episode 3

…where I talk about what should be coming

Ennustaminen on vaikeaa, erityisesti tulevaisuuden ennustaminen.
– one of my Dad's favorite sayings

(Making predictions is difficult, especially making predictions about the future.)

# Things I see in the future

I am really excited about the future

We can finally see the full realization of the original vision

…but I also see this:

1. People unsure how to build agents
2. People having unrealistic expectations about Generative AI
3. People confused about “RDF vs. Labeled Property Graphs”
4. …

# Agents!

The Semantic Web vision is predicated on the idea that we can converse with our agents and give them tasks to perform

Using LLMs, sufficiently flexible and open-ended conversational user interfaces are finally possible

Through curated and audited knowledge graphs, we get trusted sources of information for the agents to consume (and avoid LLM hallucinations)

# Agents!

The realization of autonomous agents minimally requires these:

1. KR & reasoning
2. Planning
3. Ability to converse with the agents

LLMs will give you #3, but not #1 or #2

- (despite what you hear)
- "agentic", "agentive", … huh?
- good news: we already have #1 and #2

# Highly trivialized flow

1. Human user asks their agent a question
2. Agent interprets #1
3. Result of #2 is translated into queries
4. Queries (from #3) are executed against a knowledge graph
5. Results from #4 are translated for the human user

# Highly trivialized flow

1. Human user gives their agent a task
2. Agent interprets #1
3. Result of #2 is translated into goals
4. Agent creates a plan to achieve the goals
5. Plan (from #4) is executed
6. Possible results from #5 are translated for the human user

# Highly trivialized flow

| | |
|---|---|
| 1. Human user gives their agent a task | (using natural language) |
| 2. Agent interprets #1 | LLM |
| 3. Result of #2 is translated into goals | KR, maybe also LLM |
| 4. Agent creates a plan to achieve the goals | planning, KR |
| 5. Plan (from #4) is executed | ? |
| 6. Possible results from #5 are translated for the human user | LLM |

# “Any sufficiently advanced technology is indistinguishable…”

Amidst all the hype around generative AI, there are beliefs that we can

- just use LLMs to construct ontologies and knowledge graphs
- use knowledge graphs to improve RAG techniques (“graphRAG”) and eliminate hallucinations

(I do not share these beliefs)

# "Any sufficiently advanced technology is indistinguishable…"

Amidst all the hype around generative AI, there are beliefs that we can

- just use LLMs to construct ontologies and knowledge graphs
- use knowledge graphs to improve RAG techniques ("graphRAG") and eliminate hallucinations

**AI winter warning**

(I do not share these beliefs)

# Can we use LLMs with the Semantic Web?

Everyone is building LLM-based systems

- some are using a (knowledge) graph for better RAG
- (this won't make hallucinations go away, just sayin'…)

Why not build Semantic Web -based systems?

- and use an LLM for the conversational UI, but not as a source of answers
- (whether you build agents or not)

(you can replace "Semantic Web" with "knowledge graph" here if you like)

# And finally… a graph is a graph is a graph, right? Well, no…

Why am I talking about this?

What does this have to do with the Semantic Web?

Labeled Property Graphs (LPGs) are gaining in popularity

People are building knowledge graphs using LPGs

- …and "re-inventing" features already in RDF – not always well
- (remember what I said about perception and adoption)

# Not “RDF vs. LPG” but rather “RDF and LPG”

We need better “alignment” between RDF and LPGs

- (making them same is not a goal)

RDF-star can be seen as contributing to this

- (caveat: not everyone sees the alignment as a useful goal)

Our (Neptune team’s) approach: “Project OneGraph”

- interoperability of query languages without translating data
- we already have openCypher-over-RDF available

# Summary

The Semantic Web is about AI: KR + agents

People thought it was about something else

The vision has now matured into modern knowledge graphs (also: I declare victory)

We still do not have agents, but can build them, and LLMs will help

People (still) trying to reinvent RDF (poorly)

## Summary

The Semantic Web is about AI: KR + agents

People thought it was about something else

The vision has now matured into modern knowledge graphs (also: I declare victory)

We still do not have agents, but can build them, and LLMs will help

People (still) trying to reinvent RDF (poorly)

## Take-aways

LLMs will help the Semantic Web, and Semantic Web technologies are absolutely needed to make LLMs succeed

Parts of the original vision have matured and become part of mainstream, now it is time to get the rest of the work done

# My heartfelt thanks to:

Tanya Shigaeva

Charles Ivie, Brian O'Keefe and the rest of the Neptune team at AWS

Juan Sequeda

Lauren Lassila

Adrian Gschwend

Members of the W3C RDF-star WG

All the people I worked with over the years

# Take-aways

LLMs will help the Semantic Web, and Semantic Web technologies are absolutely needed to make LLMs succeed

Parts of the original vision have matured and become part of mainstream, now it is time to get the rest of the work done

# Thank you! Any questions?

Contact: ora@amazon.com


## Take-aways

LLMs will help the Semantic Web, and Semantic Web technologies are absolutely needed to make LLMs succeed

Parts of the original vision have matured and become part of mainstream, now it is time to get the rest of the work done