Authors: Teodora Petkova, Atanas Kiryakov

The Semantic Web, both as a research field and a technology stack, is seeing mainstream industry interest, especially with the knowledge graph concept emerging as a pillar for data well and efficiently managed. But what exactly are we talking about when we talk about the Semantic Web? And what are the commercial implications of semantic technologies for enterprise data?

The Semantic Web started in the late 90's as a fascinating vision for a web of data, which is easy to interpret by both humans and machines. One of its pillars are ontologies that represent explicit formal conceptual models, used to describe semantically both unstructured content and databases. While Semantic Web is often condemned for being too academic, two of its incarnations already enjoy massive adoption.

The first one is Schema.org: millions of web pages are tagged with semantic annotations to enable a much better web search experience. The second one is the Linked Open Data (LOD): a cloud of interlinked structured datasets published without centralized control across thousands of servers. Knowledge graphs (KG) came later, but quickly became a powerful driver for the adoption of Semantic Web standards and all species of semantic technology implementing them. KGs bring the Semantic Web paradigm to the enterprises, by introducing semantic metadata to drive data management and content management to new levels of efficiency and breaking silos to let them synergize with various forms of knowledge management. This way KGs help organizations smarten up proprietary information by using global knowledge as context for interpretation and source for enrichment.

In this post you will discover the aspects of the Semantic Web that are key to enterprise data, knowledge and content management. We will walk you through the Semantic Web roots, the debate about it and further take you beyond its academic and visionary aspect into the world of efficient enterprise data management with semantic technologies and knowledge graphs. Connecting the dots for you between concepts like RDF, semantic annotation, Linked Open Data, we will help you understand why the Semantic Web will always work and how knowledge-intensive domains and applications can benefit from its affordances.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

Source: tag.ontotext.com

> In 1994 Tim Berners Lee described the Web as a "a flat, boring world devoid of meaning" (Plenary Talk Geneva) for computers. To set the stage for a Web more meaningful to our machines, in 1998, a Semantic Web Road Map came into being. After the design of this "architectural plan untested by anything except thought experiments", in 2001 the Semantic Web seized the publics' imagination with a seminal article featured in Scientific American called: The Semantic Web A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities. Today, 20 years after Tim Berners-Lee, James Hendler and Ora Lassila outlined a Semantic Web driven world where intelligent software agents automatically book flights and hotels and give us personalized answers, we can see Google transitioning from a search engine into a question-answering system, and Alexa becoming a device that can book a flight for you. Schema.org marked up pages are more than 80 million, and the LOD cloud currently contains 1,255 datasets with 16,174 links (as of May 2020).

In 1994 Tim Berners Lee described the Web as “a flat, boring world devoid of meaning” (Plenary Talk Geneva) for computers. To set the stage for a Web more meaningful to our machines, in 1998, a Semantic Web Road Map came into being. After the design of this “architectural plan untested by anything except thought experiments”, in 2001 the Semantic Web seized the public’s imagination with a seminal article featured in Scientific American called: The Semantic Web: A new form of Web content that is meaningful to computers will unleash a revolution of new possibilities.

Today, 20 years after Tim Berners-Lee, James Hendler and Ora Lassila outlined a Semantic Web driven world where intelligent software agents automatically book flights and hotels and give us personalized answers, we can see Google transitioning from a search engine into a question-answering system and Alexa becoming a device that can book a flight for you. There are more than 80 million pages with semantic, machine interpretable metadata, according to

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

the Schema.org standard. Take this restaurant, for example. Under the hood of its web content lie formalizations describing its address, opening hours, name and other details.

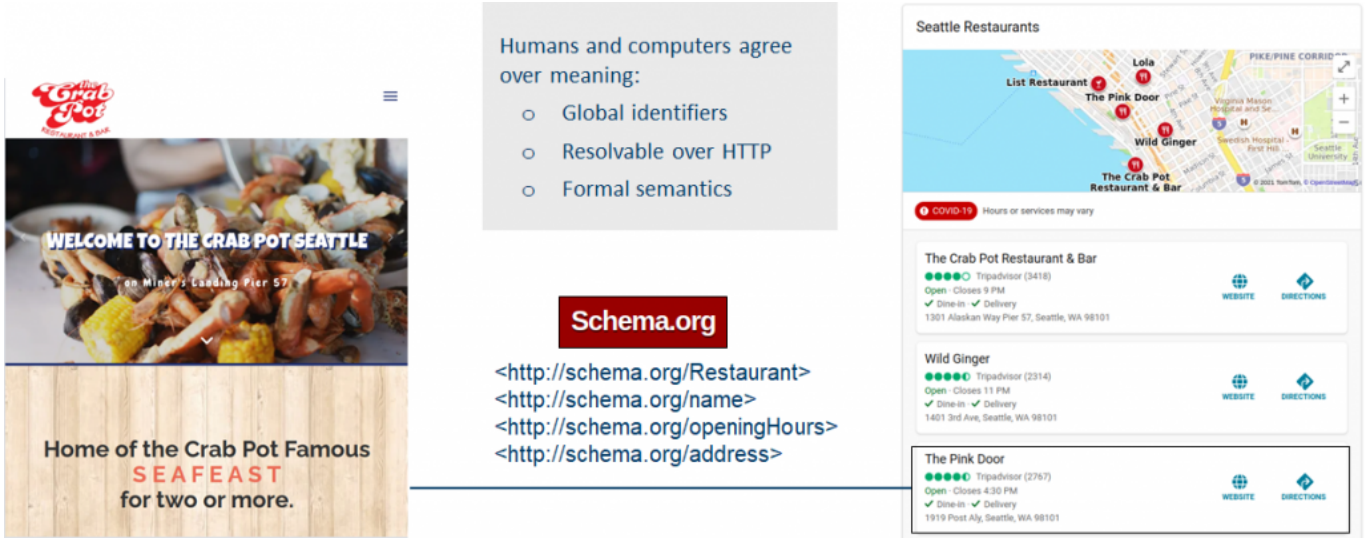

In parallel, we have the Linked Open Data (LOD) cloud, which currently contains 1,301 datasets with 16,283 links – datasets with mappings across them (https://lod-cloud.net/ as of May 2021).

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

The Semantic Web: 20 Years And a Handful of Enterprise Knowledge Graphs Later

Authors: Teodora Petkova, Atanas Kiryakov

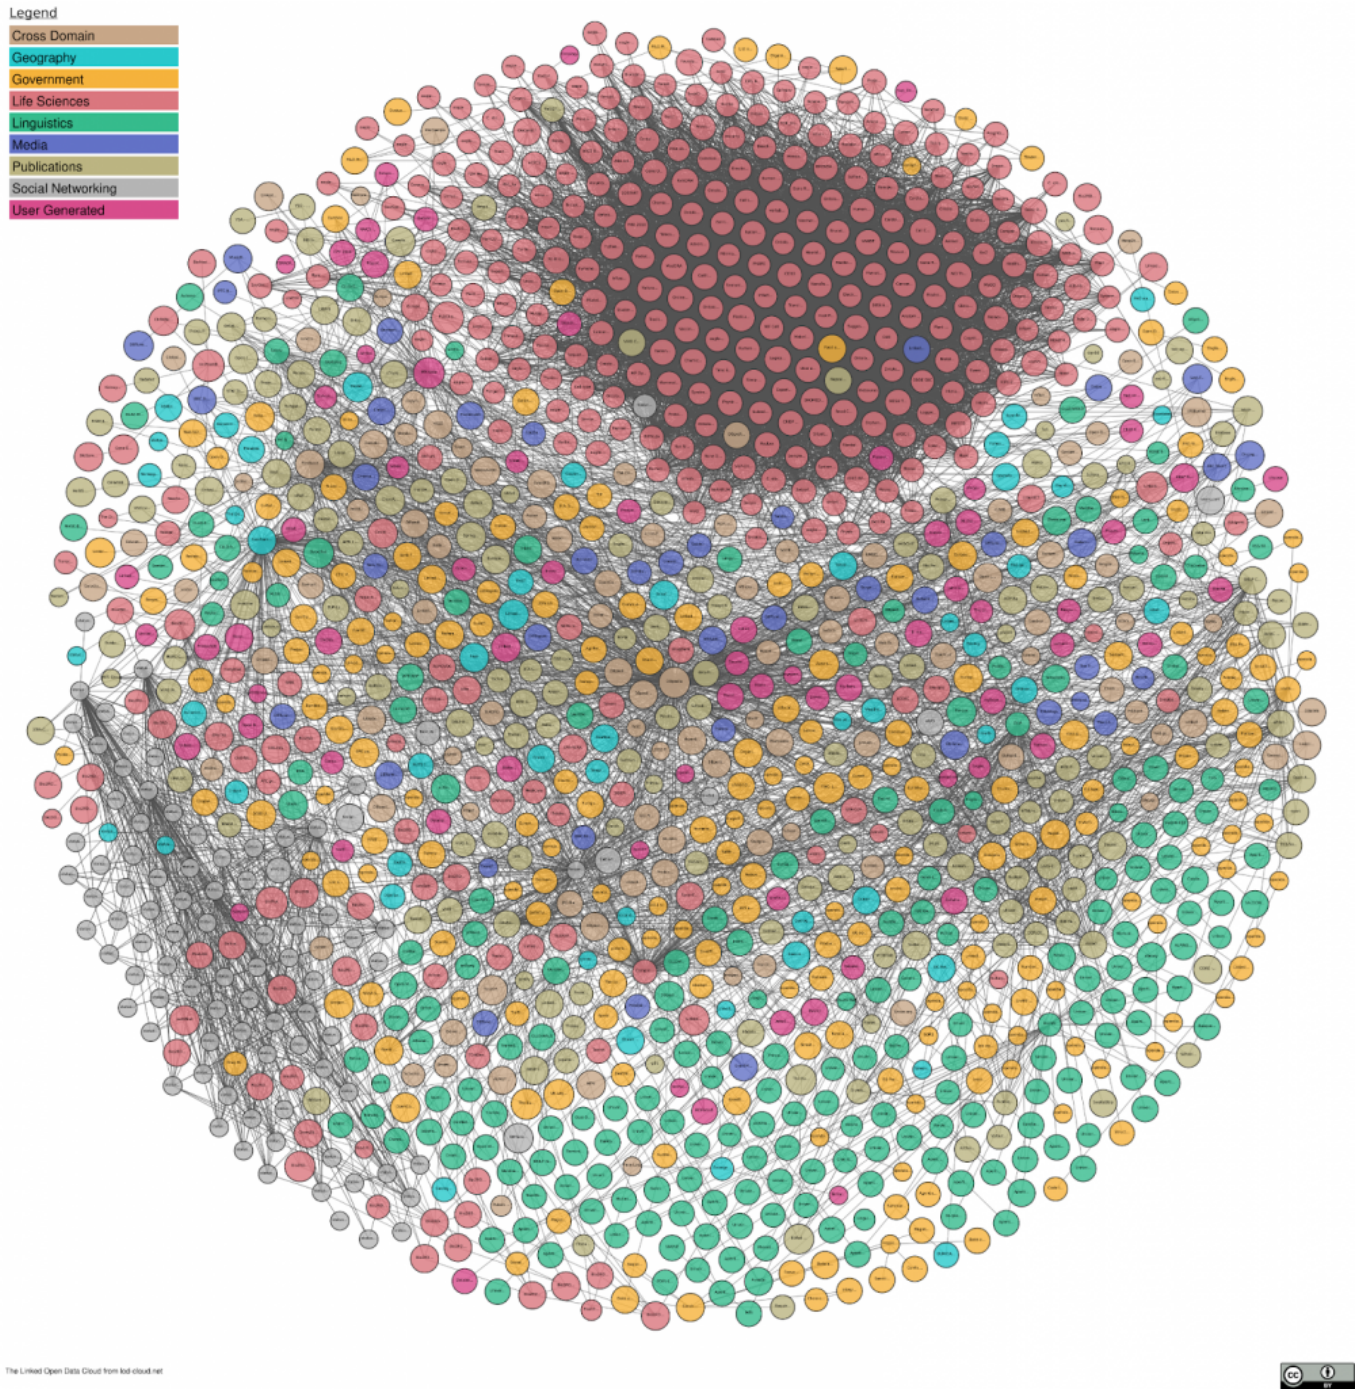

Schema.org and Linked Open Data are just two incarnations of the Semantic Web vision.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

Those are two massive information domains using W3C's stack of the Semantic Web standards (RDF, SPARQL, OWL, etc.).

And still there is confusion around the Semantic Web.

What is it? How exactly is the Web semantic? What can it do and how are enterprise knowledge graphs related to it?

Let's start with the first question.

Which Semantic Web?

Paradoxically, the Semantic Web, which aims to unambiguously describe things, people and concepts, is itself an ambiguous term.

For some, the Semantic Web is about intelligent agents browsing the Web and executing sophisticated tasks. For others, the concept boils down to smart and efficient data management (Pascal Hiltzer traced the triumphs and challenges of two decades of Semantic Web research and applications in a recent Review of the Semantic Web Field). For its critics, it is a pipe-dream, too academic to be realized, a never to come true vision (see the perspectives regarding the degree to which the original Semantic Web vision has been realised and the impact it can potentially have on the Web gathered by Aidan Hogan.) Authors Marshal and Shipman, for example, explored the many rhetorical, theoretical and pragmatic perspectives on the Semantic Web, discerning three major threads in them. The Semantic Web, their research showed, is seen as:

> (1) a universal library, to be readily accessed and used by humans in a variety of information use contexts; (2) the backdrop for the work of computational agents completing sophisticated activities on behalf of their human counterparts; and (3) a method for federating particular knowledge bases and databases to perform anticipated tasks for humans and their agents.

17 years later, in the third edition of *The Semantic Web for the Working Ontologist*, authors Dean Allemang, Hedler and Gandon covered these same perspectives with one distilled

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

explanation:

> The Semantic Web faces the problem of distributed data head-on.

In more detail, they explained that just as the hypertext Web changed how we think about the availability of documents, the Semantic Web is a radical way of thinking about data. Its main idea is to support a distributed Web at the level of the data where organizations or individuals don't just publish a human-readable presentation of information but a distributable, machine-readable description of the data.

We can see all of the above theoretical investigations in practice when it comes to what Tim Berners Lee said once when asked about the Semantic Web: a way to pull data and then pull other data and connect them to see how things fit in.

What Does The Semantic Web Already Do for Us?

If you've used Google, you've used the cornucopia of Linked data across the Web, through Google's Knowledge Graph (Google's Knowledge Graph is reportedly supported by Freebase – the knowledge acquired by Google in 2010.) If you've enjoyed the efficiency of rich snippets, you've enjoyed the riches schema.org (based on RDF) brings to the world of search since 2011. If you've used Wikidata – the structured encyclopedia – you've been using a giant RDF knowledge graph, describing about 100 million topics with over 10 billion properties and relationships. That is also one of the sources from which Google's Knowledge Graph is updated.

Taking a closer look at these applications, we see two main perspectives from which the Web is becoming increasingly semantic.

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com

Authors: Teodora Petkova, Atanas Kiryakov

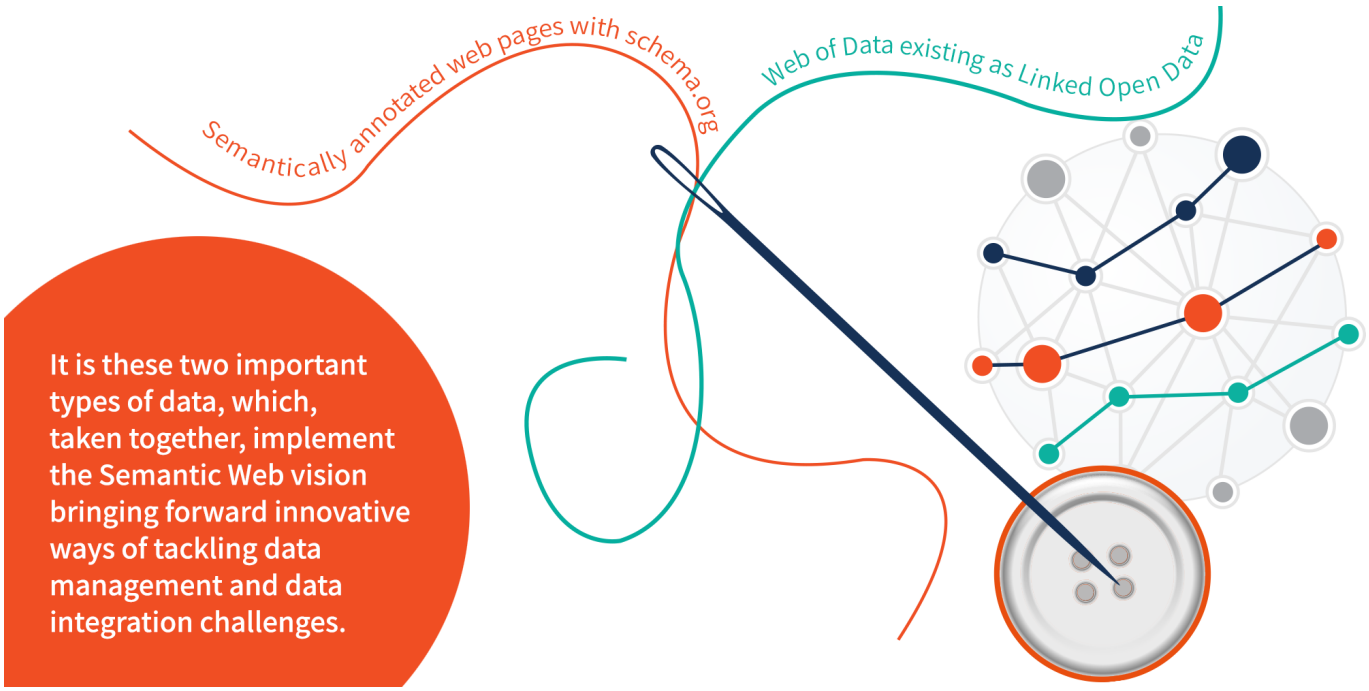

Weaving the Semantic Web with Semantic Annotations and Linked Open Data

Ontotext was founded in 2000 with the Semantic Web in its genes and we had the chance to be part of the community of its pioneers. We can’t imagine looking at the Semantic Web as an artifact. We rather see it as a new paradigm that is revolutionizing enterprise data integration and

Copyright: Ontotext New York | Sofia | Basel
Reach out to us: info@ontotext.com