




Introduction to Knowledge Technologies

Knowledge Graphs, the Semantic Web and Linked Data

Slides by Manolis Koubarakis, adapted by Isambo Karali

The Course's Content

- Introduction to the Knowledge Graphs, the Semantic Web and Linked Data
- The Resource Description Framework (RDF, RDFS)
- The query language SPARQL and its formalization
- Description logics, tableaux techniques
- The Web Ontology Language (OWL)
- Ontology Engineering
- Rule languages for the Semantic Web
- Linked geospatial data and geospatial knowledge graphs

Presentation Outline

- Discussion about the course's title and the concept of “knowledge”
- Knowledge Graphs, the vision of the Semantic Web and Linked Data
- Uses in today's search engines e.g., by Google
- RDF, RDFS and SPARQL
- Description logics, ontologies and OWL, very large ontologies
- Linked geospatial and temporal data
- Applications

The Course's Title

- **Technologies** for **Knowledge** Representation
- **Knowledge**: no clear definition, e.g. “Knowledge can be defined as the body of facts and principles accumulated by humankind or the act, fact, state of knowing” [D. Patterson, "Introduction to Artificial Intelligence and Expert Systems“, 1990]
- Requirement for human intelligence
- Uses data and relationships between them
- Human intelligence exhibits the ability to generate new knowledge from the existent one (i.e. reason about it)

- Knowledge Representation (and Reasoning): subfield of **Artificial Intelligence**
- **Nowadays:** Knowledge representation **technologies** have been developed, inspired from the ones of the world wide web, to materialize Tim **Berners-Lee**'s vision for a **Semantic Web**
- Knowledge is the key concept for modern search facilities



A bit of historical background: Semantic Networks

- A semantic network is a graph structure for representing knowledge in patterns of interconnected nodes and arcs
- Nodes (drawn as circles or boxes) represent conceptual units
- Arcs are directed links (drawn as arrows between the nodes) representing the relations between the units
- Assign a meaning to each node and link (then, becomes “semantic”)
- Special relationship represents taxonomic hierarchy (“is a” relationship)
- Some specializations restrict relationships to allow formalization by simpler algebras
- Back to ancient years (Greek philosophy)

- “Computer implementations of semantic networks were first developed for artificial intelligence and machine translation, but earlier versions have long been used in philosophy, psychology, and linguistics” [John F. Sowa, “Semantic Networks”, 1987]
- Notable implementations: Robert F. Simmons (1963), M. Ross Quillian (1963)
- Notable specialization: Sowa, John F. (1976) : Conceptual Graphs

- [Fritz Lehmann, “SEMANTIC NETWORKS”, 1992]


Toby the hungry tiger follows his mother.
 Toby, the son of the tigress he follows, is hungry.
 The tiger is followed by her hungry cub Toby.
 Hunger grips Toby, son of the tigress leading him.
 She who bore hungry tiger Toby
 is also by him followed.
 Der Tiger Toby, der seiner
 Mutter folgt, hat Hunger.
 Hungry Toby's relation to the tiger
 is one of mother-following.
 $\exists(\text{Toby})\exists(x)(\text{Tiger}(\text{Toby}) \wedge$
 $\text{Tiger}(x) \wedge \text{Follows}(\text{Toby}, x) \wedge$
 $\text{Mother-of}(x, \text{Toby}) \wedge \text{Female}(x) \wedge$
 $\text{Male}(\text{Toby}) \wedge \text{Hungry}(\text{Toby}))$.



A bit of historical background: Frames

- We put together all the knowledge about a concept (or entity) into a single structure (frame)
- Introduced by Marvin Minsky (1974)
- “A *frame* is a data-structure for representing a stereotyped situation, like being in a certain kind of living room, or going to a child's birthday party. Attached to each frame are several kinds of information. Some of this information is about how to use the frame. Some is about what one can expect to happen next. Some is about what to do if these expectations are not confirmed.” [M. Minsky, “A Framework for Representing Knowledge”, MIT-AI Laboratory Memo 306, June, 1974]

Representation of the situation of “looking at a cube” [M. Minsky, “A Framework for Representing Knowledge”, MIT-AI Laboratory Memo 306, June, 1974]



A bit of historical background: The System KL-ONE

- Ronald J. Brachman, James G. Schmolze , 1985
- “KL-ONE is a system for representing knowledge in Artificial Intelligence programs, more or less in the tradition of semantic networks and frames” [Ronald J.Brachman , James G.Schmolze, “An Overview of the KL-ONE Knowledge Representation System”, 1985]
- First appeared in 1977 in Brachman’s PhD

- Conceptual information based on the idea of structured inheritance networks
- Object centered language
- Structure types: “Concept” , “Role”
- Structure-forming operations: “specialization,” “restriction”, “differentiation”
- Whether one description “subsumes” another can also be answered by looking up a hierarchy
- “Classification” etc.

[Ronald J.Brachman , James G.Schmolze, “An Overview of the KL-ONE Knowledge Representation System”, 1985]





Figure 1. A simple KL-ONE network of Generic Concepts.

[Ronald J.Brachman , James G.Schmolze, “An Overview of the KL-ONE Knowledge Representation System”, 1985]



Towards the present: The Web

- **Universal Resource Locators (URLs)** to identify documents
- The **Hypertext Transfer Protocol (HTTP)** to exchange documents between a client and a server
- **HTML** for marking up information to be presented to human readers through a **browser**
- **Search Engines (Google!)** to discover information

Using Search Engines – Example I

- Assume that you want to find information about the actor Sean Connery.
- What do you do?
 - Google “Sean Connery”!

Result (2019)



[All](#) [Images](#) [News](#) [Videos](#) [Maps](#) [More](#) [Settings](#) [Tools](#)

About 11,300,000 results (0.64 seconds)

Sean Connery - Wikipedia

https://en.wikipedia.org/wiki/Sean_Connery ▼

Sir Thomas Sean Connery is a retired Scottish actor and producer who has won an Academy Award, two BAFTA Awards (one of them being a BAFTA Academy ...

[Sean Connery filmography](#) · [Jason Connery](#) · [Diane Cilento](#) · [Neil Connery](#)

Sean Connery - IMDb

www.imdb.com/name/nm0000125/ ▼

Sean Connery is best known for portraying the character James Bond, starring in seven Bond films between 1962 and 1983. In 1988, Connery won the Academy Award for Best Supporting Actor for his role in The Untouchables (1987). ... James Bond (voice, as Sir Sean Connery)

Sean Connery - Biography - IMDb

www.imdb.com/name/nm0000125/bio ▼

Sean Connery is best known for portraying the character James Bond, starring in seven Bond films between 1962 and 1983. In 1988, Connery won the ...

Birth Name: Thomas Sean Connery **Nickname:** Big Tam
Born: August 25, 1930 in Edinburgh, Scotland, ... **Height:** 6' 2" (1.88 m)

The Official Website of Sir Sean Connery

www.seanconnery.com/ ▼

Retrieving Messages: 'Dear Sean, I am a proud Scots Australian and salute your work for Scots independence! 'The Man who would be King' was a film beloved ...

Sean Connery - Actor, Producer - Biography.com

<https://www.biography.com/people/sean-connery-9255144> ▼

Apr 27, 2017 - Read about the childhood, career and personal life of Academy Award-winning Scottish actor Sean Connery, known as spy 007 in the early ...

Birth Date: August 25, 1930 (age 86) **Nicknames:** Big Tam

[More images](#)

Sean Connery

Actor

seanconnery.com

Sir Thomas Sean Connery is a retired Scottish actor and producer who has won an Academy Award, two BAFTA Awards and three Golden Globes. [Wikipedia](#)

Born: August 25, 1930 (age 87), Fountainbridge
Height: 1.89 m
Spouse: [Micheline Roquebrune](#) (m. 1975), [Diane Cilento](#) (m. 1962–1973)
Children: [Jason Connery](#)

Quotes [View 2+ more](#)

There is nothing like a challenge to bring out the best in man.

I like women. I don't understand them, but I like them.

Result (today)



sean connery

All

Images

News

Videos

Short videos

Forums

Web

More

Tools

Sean Connery

Actor

Overview

Movies

Songs



Sir Sean Connery in Rome for the Rome Film Festival where he will be presented with a...





007.com

Sean Connery: Six Great Moments | James Bond 007

Born on 25 August 1930, Sean Connery became the first James Bond on film in... 1 month ago

Born

Aug 25, 1930

Edinburgh
Royal Matern...

Died

Oct 31, 2020

Lyford Cay

See also



Dr. No

1962 film



James Bond

Film character

Movies



Dr. No

1962



Goldfinger

1964



The Rock

1996



Indiana Jones and the Last Crusade

1989



From Russia with Love

1963



Never Say Never Again

1983

Wikipedia

https://en.wikipedia.org/wiki/Sean_Connery

Sean Connery

He was the first actor to portray the fictional British secret agent James Bond in motion pictures, starring in seven Bond films between 1962 and 1983.

Sean Connery filmography

Jason Connery

Neil Connery

Diane Cilento

About

Sir Thomas Sean Connery was a Scottish actor. He was the first actor to portray the fictional British secret agent James Bond in motion pictures, starring in seven Bond films between 1962 and 1983. [Wikipedia](#)

Born: August 25, 1930, [Edinburgh Royal Maternity Hospital and Simpson Memorial Maternity Pavilion, United Kingdom](#)

Died: October 31, 2020 (age 90 years), [Lyford Cay](#)

Spouse: [Micheline Roquebrune](#) (m. 1975–2020), [Diane Cilento](#) (m. 1962–1974)

Children: [Jason Connery](#)

Siblings: [Neil Connery](#)

Height: 1.88 m

Feedback

People also search for



Roger Moore



Micheline



Jason



Pierce

Discussion

- Gone are the days when Google would simply return a list of links to Web pages containing the words in the query ordered by importance (using PageRank)
- Now Google is able to understand that a user question is about an **entity** and returns **structured information about this entity** (the knowledge panel) together with relevant web links

Using Search Engines – Example II

- Assume you would like to find information about Prof. Manolis Koubarakis.
- Google “Prof. Manolis Koubarakis”.

Result



[All](#) [Images](#) [News](#) [Videos](#) [Short videos](#) [Web](#) [Books](#) [More](#) [Tools](#)

**Εθνικόν και Καποδιστριακόν Πανεπιστήμιον Αθηνών**
<https://cgi.di.uoa.gr/~koubarak/>

**Manolis Koubarakis**
Professor, Department of Informatics and Telecommunications, National and Kapodistrian University of Athens, University Campus, Ilisia 15784 Athens Greece.

**Google Scholar**
<https://scholar.google.com/citations>

Manolis Koubarakis
Professor, Dept. of Informatics and Telecommunications, National and Kapodistrian University of -
Παράτιςται από 9.165 - Artificial Intelligence ...

**MaDgIK**
<https://www.madgik.di.uoa.gr/~m...> [Translate this page](#)

Μανόλης Κουμπαρακής - Αθήνα
Ο Μανόλης Κουμπαρακής είναι Καθηγητής του Τμήματος Πληροφορικής και Τηλεπικοινωνιών του Πανεπιστημίου Αθηνών από τον Νοέμβριο του 2005.

**LinkedIn · Manolis Koubarakis**
1.9K+ followers

Manolis Koubarakis - Professor at Dept.of Informatics and ...
Greece · Professor · Dept.of Informatics and Telecommunications, National and Kapodistrian University of Athens
My research concentrates on Artificial Intelligence and Big Data. See my Web page <http://cgi.di.uoa.gr/~koubarak/> and the Web site of my team
Missing: **Prof.** | Show results with: **Prof.**

Images

**Manolis Koubarakis**
cgi.di.uoa.gr

**mkoubarakis | MaDgIK**
MaDgIK

**Manolis Koubarakis**
Archimedes Research Unit

Discussion

- Good work but no structured information about this entity.

Using Search Engines – Example III

- Let us now google: “Who is married to Angelina Jolie?”

Result

File Edit View History Bookmarks Profiles Tools Help

who is married to angelina jolie

who is married to angelina jolie

Google

who is married to angelina jolie

AI Mode All Images News Videos Forums Short videos More Tools

Angelina Jolie / Spouse

 Brad Pitt m. 2014–2019	 Billy Bob Thornton m. 2000–2003	 Jonny Lee Miller m. 1996–1999
--	---	---

AI Overview

Angelina Jolie is currently single and **does not have a husband**.

She has been married three times in the past:

Former Husbands

- Jonny Lee Miller:** Married from 1996 to 1999. They met while co-starring in the film *Hackers*.
- Billy Bob Thornton:** Married from 2000 to 2003. Their intense, highly publicized relationship included wearing vials of each other's blood.
- Brad Pitt:** Married from 2014 to 2019, with their divorce finalized in late 2024. They

Show more

Facebook

Angelina Jolie: three marriages — three completely different versions...

Three marriages. Three men. Three very different eras of Angelina Jolie. Her first...

People.com

Angelina Jolie's Dating History: From Billy Bob Thornton to Brad Pitt

Spouses and Partners: Jonny Lee Miller (married 1996, divorced 1999), Jonny Shimiz...

Wikipedia

https://en.wikipedia.org/wiki/Angelina_Jolie

Angelina Jolie

Spouses: Jonny Lee Miller (m. 1996; div. 2000), Billy Bob Thornton (m. 2000; div. 2003), Brad Pitt (m. 2014; div. 2024)

2

Windows Search Task View Microsoft Edge Google Chrome Brave Firefox

ENG 9:18 AM 10/2/2026

Discussion

- Not only Google knows about entities, it also knows about **relationships** e.g., “is married to”.
- Google can understand the **semantics (meaning)** of a user query.

The Google Knowledge Graph

The screenshot displays the Google Inside Search Knowledge Graph interface. The top navigation bar includes links for Home, How Search Works, Tips & Tricks, Features, Search Stories, Playground, Blog, and Help. The main visual is a dark space-themed background with a network of nodes and connecting lines. A large, prominent blue circular node features a portrait of Leonardo da Vinci. Other smaller nodes include a Canadian flag, a globe, and various historical figures. On the right side, a detailed knowledge panel for Leonardo da Vinci is visible, containing a portrait, a biographical summary, and key dates. Above this panel, a carousel of four small portrait images is shown. At the bottom of the interface, there are two call-to-action buttons: a blue one labeled 'The Knowledge Graph' and a grey one labeled 'See it in action', both with right-pointing arrows. The text 'Learn more about one of the key breakthroughs behind the future of search.' is positioned below the blue button, and 'Discover answers to questions you never thought to ask, and explore collections and lists.' is below the grey button. A decorative line graph with nodes is visible at the very bottom of the page.

Google Inside Search

Home How Search Works Tips & Tricks Features Search Stories Playground Blog Help

The Knowledge Graph
Learn more about one of the key breakthroughs behind the future of search.

See it in action
Discover answers to questions you never thought to ask, and explore collections and lists.

Leonardo da Vinci
Leonardo di ser Piero da Vinci was an Italian Renaissance polymath: painter, sculptor, architect, musician, scientist, mathematician, engineer, inventor, anatomist, geologist, cartographer, botanist, and writer.

Born: April 15, 1452, Anchiano
Died: May 2, 1519, Clos Lucé
Buried: Château d'Amboise
Parents: Caterina da Vinci, Piero da Vinci
Structures: Vitruvian Man, Da Vinci Project

Related images:
Lady with an Ermine... 1489
Ginevra de' Benci 1478
The Virgin a... 1508
Adoration of the M... 1481

Discussion

- One important technology that drives Google search is the **Google Knowledge Graph**
- The Google Knowledge Graph stores information about real-world entities and relationships
- Some knowledge graphs are used primarily within the organisation that created them (e.g. Google knowledge graph, Amazon's product graph)
- Other knowledge graphs are openly available. These include **DBpedia, Wikidata, WordNet, Geonames**, etc.

- The emphasis has shifted to **knowledge graphs**. See, for example, the proceedings of a Dagstuhl seminar on the topic:
http://drops.dagstuhl.de/opus/volltexte/2019/10328/pdf/dagrep_v008_i009_p029_18371.pdf
- Knowledge graphs is a special case of Knowledge bases
- Constructing and using knowledge bases has been, for a long time, the subject of research in Knowledge Representation and Reasoning
- The problem of knowledge representation and reasoning on the Web has been studied by the areas of Semantic Web and Linked Data, which are the topic of this course

Using Search Engines – Example VI

- More “difficult questions”:
“why HDL is no longer taken into account as a risk avoidance factor?”

Result



why HDL is no longer taken into account as a risk avoidance factor?



All Short videos Forums Images Videos Web News More ▾ Tools ▾

◆ AI Overview

HDL is no longer the sole focus in cardiovascular risk assessment because simply increasing HDL levels through medication doesn't reduce heart attack risk, the complexity of HDL particles, and the potential for very high HDL to be harmful are important factors. Instead, current guidelines emphasize the importance of targeting «good» non-HDL cholesterol and overall triglyceride levels, alongside established factors like LDL cholesterol, blood pressure, and lifestyle changes. ⓘ

Why the shift away from HDL?

- **Lack of successful therapies:** Clinical trials using drugs designed to raise HDL have failed to show a reduction in cardiovascular events like heart attacks. ⓘ
- **Complexity of HDL particles:** HDL isn't a single entity; it's a family of different particles with varying sizes and functions. A high HDL level doesn't guarantee a functional, protective particle. ⓘ
- **Risk of paradoxically high levels:** Studies indicate that extremely high levels of HDL cholesterol can be associated with increased mortality, suggesting a potential for harm rather than benefit. ⓘ
- **HDL as a marker, not a cause:** The evidence suggests that HDL might be a biomarker that reflects overall health rather than being a direct cause of disease



No-HDL-cholesterol as risk marker and therapeutic goal

In fact, according to clinical practice guideline recommendations and the results from various clinical...

Elsevier ⓘ

The complex story of HDL cholesterol - Harvard Health

Aug 1, 2025

Harvard Health ⓘ



Lipids Series-HDL Can It Be Too High? Is Treatment Available?

Jan 4, 2021 — here the hazard ratio being closer to unity. so these data suggests th...

YouTube · Mayo Clinic ⓘ



- Google uses generative AI in Search
- Overviews generated with the help of AI
- Info from across the web and Google's Knowledge Graph
- Work still in progress (feedback needed by users)

LLMs and KGs

- LLMs and KGs are not opposing forces, they are combined in various themes:
 - LLMs for KGs:
 - Knowledge Extraction and Canonicalisation (entity resolution, knowledge extraction)
 - KG Construction (link prediction, triple extraction)
 - Ontological Schema Construction
 - KGs for LLMs:
 - Training and Augmenting LLMs

Towards a Better Web!

- The **Semantic Web** vision was articulated in a Scientific American article by Tim Berners-Lee, James Hendler and Ora Lassila (May 2001).
 - “The Semantic Web will bring **structure to the meaningful content of Web pages**, creating an environment where **agents** roaming from page to page readily **carry out sophisticated tasks for users**.”
- You can find the article on various Web sites e.g., <http://www.dcc.uchile.cl/~cgutierrez/cursos/IC/semantic-web.pdf>
- Notice the words **meaning** (semantics) and **agents** (a role for all of Artificial Intelligence).

Original Introduction of the Term

Berners-Lee, Tim.

Weaving the Web : the original design and ultimate destiny of the World Wide Web by its inventor / Tim Berners-Lee with Mark Fischetti. — 1st ed.

Includes index.

ISBN 0-06-251586-1 (cloth)

ISBN 0-06-251587-X (paper)

1. World Wide Web (Information retrieval system)—History. 2. Berners-Lee, Tim. I. Fischetti, Mark. II. Title.

5105.888 .B46 1999

5.04—dc21

99-27665

I have a dream for the Web . . . and it has two parts.

In the first part, the Web becomes a much more powerful means for collaboration between people. I have always imagined the information space as something to which everyone has immediate and intuitive access, and not just to browse, but to create. The initial *WorldWideWeb* program opened with an almost blank page, ready for the jottings of the user. Robert Cailliau and I had a great time with it, not because we were looking at a lot of stuff, but because we were writing and sharing our ideas. Furthermore, the dream of people-to-people communication through shared knowledge must be possible for groups of all sizes, interacting electronically with as much ease as they do now in person.


In the second part of the dream, collaborations extend to computers. Machines become capable of analyzing all the data on the Web—the content, links, and transactions between people and computers. A “Semantic Web,” which should make this possible,

has yet to emerge, but when it does, the day-to-day mechanisms of trade, bureaucracy, and our daily lives will be handled by machines talking to machines, leaving humans to provide the inspiration and intuition. The intelligent "agents" people have touted for ages will finally materialize. This machine-understandable Web will come about through the implementation of a series of technical advances and social agreements that are now beginning (and which I describe in the next chapter).

Once the two-part dream is reached, the Web will be a place where the whim of a human being and the reasoning of a machine coexist in an ideal, powerful mixture.

It will require a lot of nitty-gritty work.

Humans in the Ladder of Causality [J. Pearl and D. Mackenzie, “The Book of Why: The New Science of Cause and Effect”, 2018]



How Can we Achieve the Semantic Web? – The Original Vision

- Instead of publishing information to be consumed by humans, publish **machine-processable data and metadata** using terms/languages that can be understood by machines.
- Build machines (**agents**) that will search for, query, integrate etc. this data.
- Make sure **all** agents understand your terms/languages.

The Semantic Web and Linked Data Vision Today

- From the Introduction section of the Semantic Web activity of the W3C
<http://www.w3.org/2001/sw/>

Linked Data

“The Semantic Web is a **web of data**. There is lots of data we all use every day, and it is not part of the web. I can see my bank statements on the web, and my photographs, and I can see my appointments in a calendar. But can I see my photos in a calendar to see what I was doing when I took them? Can I see bank statement lines in a calendar?


Why not? Because we don't have a web of data. Because data is controlled by applications, and each application keeps it to itself.

The Semantic Web is about two things. It is about **common formats for integration and combination of data drawn from diverse sources**, where on the original Web mainly concentrated on the interchange of documents. It is also about **language for recording how the data relates to real world objects**. That allows a person, or a machine, to start off in one database, and then move through an unending set of databases which are connected not by wires but by being about the same thing.”

The Semantic Web Vision Today (cont'd)

- Stressing the growing need for **linked data** and **Web data integration**. Lots of important Web applications demand this e.g., **e-science** and **e-government**.
- Stressing the need for **Web standards/languages for expressing shared meaning**. This is important if we want **agents** that are not handcrafted only for particular tasks to be developed.
- See the paper “The Semantic Web Revisited” by Nigel Shadbolt, Wendy Hall and Tim Berners-Lee at http://eprints.ecs.soton.ac.uk/12614/1/Semantic_Web_Revisited.pdf .
- Agentic AI (now) VS Multi-Agent Systems (late '80s) ← Distributed Artificial Intelligence (DAI) ('70s)


The Semantic Web Layer Cake



Emphasis of this Course

- The Semantic Web topics that we will cover in this course are:
 - Linked data
 - RDF, RDFS, SPARQL and RDF stores (Sesame).
 - Ontologies
 - Description logics, OWL 2, tools for developing ontologies (Protégé) and reasoners (Pellet, Hermit).
 - Rules
 - SWRL and others
 - Ontology engineering
 - Linked geospatial data (the work of prof. Koubarakis' group)
 - stSPARQL, GeoSPARQL and Strabon

Outline

- The Semantic Web vision
- **Linked data** 
- Very large knowledge bases
- Ontologies
- Linked geospatial and temporal data
- Applications

Open Data

- The open data movement is about **making data public** so they can be used freely by everyone.
- Lots of such data becoming available today is **government data**.



<http://data.gov>

File Edit View History Bookmarks Profiles Tools Help

Data.gov Home - Data.gov

data.gov

An official website of the United States government [Here's how you know](#)

Try the next-generation Data Catalog at catalog.data.gov and help shape it with your feedback.

 DATA.GOV

DATA METRICS OPEN GOVERNMENT CONTACT

 [User Guide](#)

The Home of the U.S. Government's Open Data

Here you will find data, tools, and resources to conduct research, develop web and mobile applications, design data visualizations, and more.

603,381 DATASETS AVAILABLE


Search



2

9:40 AM 10/2/2026

<http://data.gov.uk>



<https://data.gov.gr/>

The screenshot shows the homepage of data.gov.gr. At the top, there's a navigation bar with the govgr logo, the text 'data.gov.gr', and language options 'Ελληνικά' and 'English'. A search bar is also present. Below the navigation bar, there's a main section with the heading 'Βρείτε ανοικτά δεδομένα' (Find open data). This section includes a search bar with the placeholder 'Αναζήτηση στο data.gov.gr' and a magnifying glass icon. To the right of the search bar, there's a list of 'Δημοφιλέστερα Σύνολα Δεδομένων' (Most Popular Data Sets), including 'Μητρώο Πολιτών', 'Αριθμός Κρατούμενων Σωφρονιστικών Καταστημάτων', 'Αριθμός Μαθητών ανά Σχολείο', and 'Μετεωρολογικά Δεδομένα Μετεωρολογικών Σταθμών Κρήτης'. Below this, there are three thematic sections: 'Οικονομία και Εμπόριο', 'Δημόσια Τάξη και Δικαιοσύνη', and 'Εκπαίδευση'. Each section has a brief description and a corresponding bar chart. The 'Οικονομία και Εμπόριο' chart shows 'Κρατικός Τομέας', 'Σύνολο χώρας', 'Τομέας Επιχειρήσεων', and 'Τομέας Ιδιωτικών Μη-Κερδοσκοπικών Ιδρυμάτων'. The 'Δημόσια Τάξη και Δικαιοσύνη' chart shows 'Εγκλήματα' and 'Αποτίσεις'. The 'Εκπαίδευση' chart shows 'Αριθμός σχολείων'.

Βρείτε ανοικτά δεδομένα

Αναζητήστε δεδομένα δημοσιευμένα από την κεντρική διοίκηση, οργανισμούς τοπικής αυτοδιοίκησης και άλλες υπηρεσίες

Αναζήτηση στο data.gov.gr

Θεματικές Ενότητες Δεδομένων

10 θεματικές ενότητες με 74 σειρές δεδομένων.

Δημοφιλέστερα Σύνολα Δεδομένων

- [Μητρώο Πολιτών](#)
Μητρώο ενεργού πληθυσμού ανά φύλο, δήμο, γεωγραφικό διαμέρισμα και έτος γέννησης
- [Αριθμός Κρατούμενων Σωφρονιστικών Καταστημάτων](#)
Στατιστικά στοιχεία επί του αριθμού των κρατούμενων σωφρονιστικών καταστημάτων προερχόμενα από τη Γενική
- [Αριθμός Μαθητών ανά Σχολείο](#)
Αριθμός μαθητών ανά σχολείο και φύλο
- [Μετεωρολογικά Δεδομένα Μετεωρολογικών Σταθμών Κρήτης](#)

Οικονομία και Εμπόριο

Οικονομία, εμπόριο, βιομηχανία, εισαγωγές, εξαγωγές

Κρατικός Τομέας, Σύνολο χώρας, Τομέας Επιχειρήσεων, Τομέας Ιδιωτικών Μη-Κερδοσκοπικών Ιδρυμάτων, Τομέας Τριτοβάθμιας και Μεταδευτεροβάθμιας Εκπαίδευσης

Δημόσια Τάξη και Δικαιοσύνη

Δικαστήρια, αστυνομία, φυλακές, εγκλήματα, σύνορα και μετανάστευση

Εγκλήματα, Αποτίσεις, Διευκολύνσεις

Εκπαίδευση

Μαθητές, φοιτητές, πανεπιστήμια, καθηγητές και βαθμολογίες


Αριθμός σχολείων

Linked Open Data

- The goal of this W3C community effort is “to **extend the Web with a data commons** by publishing various open data sets as RDF on the Web and by setting RDF links between data items from different data sources.”
- The LOD community is developing a set of **best practices** for achieving this.



The Linked Open Data Cloud



See <http://lod-cloud.net/> for more details

The Greek Administrative Geography (GAG)



- According to the Kallikratis administrative reform of 2010, Greece consists of:
 - 7 **decentralized administrations** (e.g., Macedonia and Thrace)
 - 13 **regions** (e.g., West Macedonia)
 - Each region consists of **regional units** (e.g., Heraklion)
 - Each regional unit consists of **municipalities** (e.g., Dimos Chersonisou)
 - ...
- You can find the dataset at <http://www.linkedopendata.gr/dataset/greek-administrative-geography>

Querying GAG

 **stSPARQL Endpoint**

On this page you can execute stSPARQL queries against the Strabon backend. The dataset is based on the Greek Administrative Geography (Kallikratis). On the left sidebar, some example stSPARQL queries are provided.

Discovery Queries

- Query 1: Find all municipalities of Peloponnese
- Query 2: Find municipalities with low population
- Query 3: Find the geometry of municipality of Athens

Explore/Modify operations

You must be logged in to perform update queries.

stSPARQL Query:

```
#Geometry of municipality of Athens
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX gag: <http://geo.linkedopendata.gr/gag/ontology/>

SELECT ?geometry
WHERE
{
  ?m rdf:type gag:Δήμος .
  ?m rdfs:label "ΔΗΜΟΣ ΑΘΗΝΑΙΩΝ" .
  ?m gag:έχει_γεωμετρία ?geometry.
}
```

Output Format:

KML

View Result:

On a map

Map Bounds:

<http://www.opengis.net/def/crs/EPSS/0/4326> POLYGON(38.06, 24.09 37.77))""<http://www.opengis.net/ont/g

Query



Key Linked Data Technologies


- **URIs**: a generic means to **identify entities or concepts in the world**
- **HTTP**: a simple, yet universal, mechanism for **retrieving resources, or descriptions of resources**
- **RDF**: a **data model** for structuring and linking data that describes things in the world
- **SPARQL**: a **query language** for querying linked data

Uniform Resource Identifiers


- The Web provides a general form of identifier, called the **Uniform Resource Identifier (URI)**, for identifying (naming) resources on the Web.
- Unlike URLs, URIs are not limited to identifying things that have network locations, or use other computer access mechanisms. A number of different **URI schemes (URI forms)** have been already been developed, and are being used, for various purposes.
- **Examples:**
 - `http:` (Hypertext Transfer Protocol, for Web pages)
 - `mailto:` (email addresses), e.g., `mailto:em@w3.org`
 - `ftp:` (File Transfer Protocol)
 - `urn:` (Uniform Resource Names, intended to be persistent location-independent resource identifiers), e.g., `urn:isbn:0-520-02356-0` (for a book)
- No one person or organization controls who makes URIs or how they can be used. While some URI schemes, such as URL's `http:`, depend on centralized systems such as DNS, other schemes, such as `freenet:`, are **completely decentralized**.

Example

- http://en.wikipedia.org/wiki/Angelina_Jolie
- http://dbpedia.org/page/Angelina_Jolie
- <http://www.imdb.com/name/nm0001401/>
- <http://myfavouriteactors.gr/AngelinaJolie.html> (does not exist right now)
- Note the difference between the **person** Angelina Jolie and the various Web pages for her. It is common to use “hash URIs” to refer to the real person:
 - http://en.wikipedia.org/wiki/Angelina_Jolie#this




Example (cont'd)



**has
director**



**has
actor**



[http://dbpedia.org/page/Changeling_\(film\)](http://dbpedia.org/page/Changeling_(film))

- [dbpedia-owl:starring](http://dbpedia.org/page/Changeling_(film)): http://dbpedia.org/page/Angelina_Jolie
- [dbpedia-owl:director](http://dbpedia.org/page/Changeling_(film)): http://dbpedia.org/page/Clint_Eastwood

Dereferenceable URIs

- Linked data design principles prescribe that URIs should be **dereferenceable**.
- **Dereferenceable URIs** allow us to **obtain a copy or representation of the resource they identify**. In this way, we can **navigate the Web** using e.g., HTTP and obtain useful information.

Problem

- Data sets will typically **talk about the same object using different URIs**. For example:



http://www.freebase.com/view/en/angelina_jolie



http://dbpedia.org/page/Angelina_Jolie

Solution: Linking

- Use the **special kind of link** `owl:sameAs` that identifies the resources:



http://www.freebase.com/view/en/angelina_jolie




http://dbpedia.org/page/Angelina_Jolie



Other Links Beyond owl:sameAs

- **Natura 2000**
 - Example of link to GAG: This protected area **is contained** in the prefecture of Macedonia.
- **Public transport**
 - Examples of links:
 - This bus stop **is in** the municipality of Athens.
 - This bus stop **is near** a public building.
- **Public buildings**

5-star Linked Open Data



https://www.w3.org/2011/gld/wiki/5_Star_Linked_Data

“Tim Berners-Lee, the inventor of the Web and initiator of the Linked Data project, suggested a 5 star deployment scheme for Linked Data. The 5 Star Linked Data system is cumulative. Each additional star presumes the data meets the criteria of the previous step(s).

- ☆ Data is available on the Web, in whatever format.
- ☆☆ Available as machine-readable structured data, (i.e., not a scanned image).
- ☆☆☆ Available in a non-proprietary format, (i.e, CSV, not Microsoft Excel).
- ☆☆☆☆ Published using open standards from the W3C (RDF and SPARQL).
- ☆☆☆☆☆ All of the above and links to other Linked Open Data. “

RDF Basics



- The world consists of **resources**.
- Resources are identified by **URIs** and described by simple **properties** and property **values**.
- **Example:** “Angelina Jolie stars in the film Changeling directed by Clint Eastwood.”

RDF Statements

- RDF statements are **triples** of the form
subject predicate object .
- **Example:**

```
<http://dbpedia.org/page/Changeling_(film)>  
  <http://dbpedia.org/ontology/starring>  
  <http://dbpedia.org/page/Angelina_Jolie> .
```

```
<http://dbpedia.org/page/Changeling_(film)>  
  < http://dbpedia.org/ontology/director>  
  <http://dbpedia.org/page/Clint_Eastwood> .
```

Example (cont'd)

- To avoid writing long URIs we can use **namespace prefixes**:

```
@prefix dbpedia: http://dbpedia.org/page/ .
```

```
@prefix dbpedia-owl: http://dbpedia.org/ontology .
```

```
dbpedia:Changeling_(film) dbpedia-owl:starring dbpedia:Angelina_Jolie .
```

```
dbpedia:Changeling_(film) dbpedia-owl:director dbpedia:Clint_Eastwood .
```

Example (cont'd)

- Triples may involve resources described in different namespaces
 - Recall the `owl:sameAs` example:


```
@prefix fbase: http://www.freebase.com/view/en/.
```

```
@prefix owl: http://www.w3.org/2002/07/owl# .
```


```
@prefix dpedia: http://dbpedia.org/page/ .
```

```
fbase:angelina_jolie owl:sameAs dpedia:Angelina_Jolie .
```

RDF Graphs




RDF Graphs (cont'd)




RDF Schema

- **RDF is W3C a standard model for data interchange on the Web.** It provides a flexible and extensible way to describe and connect entities, enabling data sharing and integration across different applications and systems.
 - Key concepts: triple, URI
 - <https://www.w3.org/TR/rdf-primer/>
- **RDFS (RDF Schema) is a W3C standard for describing RDF vocabularies and data models.** It provides a set of classes, properties, and utility properties for defining ontologies, enabling richer integration and interoperability of data among descriptive communities.
 - Key concepts: class, property, subclass, subproperty
 - <https://www.w3.org/TR/rdf-schema/>
- The RDFS is like a **schema definition language** in the relational or object-oriented data models.

Example – Class Hierarchy



Example – Property Hierarchy



Standard Vocabularies

- Linked data are written using standard vocabularies i.e., **vocabularies that have been agreed by certain communities for describing certain kinds of resources.**
- Examples:
 - FOAF
 - Dublin Core
 - schema.org

FOAF



- The Friend of a Friend (FOAF) project has created a standard vocabulary for **describing people, the links between them and the things they create and do.**
- See <https://xmlns.com/foaf/spec/> for more information.

Example


```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .

<http://en.wikipedia.org/wiki/Angelina_Jolie#this>
  rdf:type foaf:Person ;
  foaf:name "Angelina Jolie" ;
  foaf:mbox <mailto:ajolie@gmail.com> ;
  foaf:homepage <http://www.AngelinaJolie.com/> ;
  foaf:interest < http://en.wikipedia.org/wiki/UNHCR> ;
  foaf:knows [ rdf:type foaf:Person ;
               foaf:name "Michael Douglas" ] .
```

Dublin Core

- Dublin Core is a community effort that has defined a set of metadata attributes (e.g., creator, creation date, title etc.) for **describing a wide range of digital resources**.
- See <http://dublincore.org> for more information.

Example



<http://schema.org>

- This was **an initiative by Bing, Google, Yahoo! and Yandex to provide a collection of schemas** that webmasters can use to markup their pages in ways recognized by major search providers
- Since April 2015, the W3C Schema.org Community group manages this activity
- As of 2024, over 45 million web domains markup their web pages with over 450 billion Schema.org objects
- <https://schema.org/version/latest/schemaorg-current-https.ttl>

<http://schema.org/Place>

Place

A Schema.org Type

Thing › Place

[\[more...\]](#)

Entities that have a somewhat fixed, physical extension.

Property	Expected Type	Description
Properties from Place		
additionalProperty	PropertyValue	A property-value pair representing an additional characteristic of the entity, e.g. a product feature or another characteristic for which there is no matching property in schema.org. Note: Publishers should be aware that applications designed to use specific schema.org properties (e.g. https://schema.org/width, https://schema.org/color, https://schema.org/gtin13, ...) will typically expect such data to be provided using those properties, rather than using the generic property/value mechanism.
address	PostalAddress or Text	Physical address of the item.
aggregateRating	AggregateRating	The overall rating, based on a collection of reviews or ratings, of the item.
	LocationFeatureSpecification	An amenity feature (e.g. a characteristic or service) of the

Microdata

- Microdata provide a simple way to **annotate HTML elements with machine readable tags**.
- Microdata can use standardized vocabularies to capture the **semantics** of HTML items.
- See <http://www.w3.org/TR/microdata/> and <https://html.spec.whatwg.org/multipage/microdata.html> for more details.
- Similar (earlier) approaches are **RDFa**, **microformats** and **JSON-LD**.

Example – Original HTML

- The following is some HTML code for describing a local business called “Beachwalk Beachwear and Giftware”.

```
<h1>Beachwalk Beachwear & Giftware</h1>
```

```
A superb collection of fine gifts and clothing to  
accent your stay in Mexico Beach.
```

```
3102 Highway 98
```

```
Mexico Beach, FL
```

```
Phone: 850-648-4200
```

Example (cont'd)– HTML With Microdata

```
<div itemscope itemtype="http://schema.org/LocalBusiness">
  <h1><span itemprop="name">Beachwalk Beachwear & Giftware</span></h1>
  <span itemprop="description"> A superb collection of fine gifts and
    clothing to accent your stay in Mexico Beach.</span>
  <div itemprop="address" itemscope
    itemtype="http://schema.org/PostalAddress">
    <span itemprop="streetAddress">3102 Highway 98</span>
    <span itemprop="addressLocality">Mexico Beach</span>,
    <span itemprop="addressRegion">FL</span>
  </div>
  Phone: <span itemprop="telephone">850-648-4200</span>
</div>
```

SPARQL

- SPARQL stands for “SPARQL Protocol and RDF Query Language”.
- It is the **standard query language for RDF data** proposed by the W3C.
- It is based on **matching graph patterns against RDF graphs**.
- The simplest kind of graph pattern is a **triple pattern**.
 - A **triple pattern** is like an RDF triple, but with the option of a **variable** in the subject, predicate or object positions.

Example Dataset

```
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .

<http://en.wikipedia.org/wiki/Angelina_Jolie#this>
  rdf:type foaf:Person ;
  foaf:name "Angelina Jolie" ;
  foaf:mbox <mailto:ajolie@gmail.com> ;
  foaf:homepage <http://www.AngelinaJolie.com/> ;
  foaf:interest <http://en.wikipedia.org/wiki/UNHCR> ;
  foaf:knows [ rdf:type foaf:Person ;
               foaf:name "Michael Douglas" ] .
```


Example SPARQL Query

```
SELECT ?name
WHERE {
  ?x foaf:name ?name .
  ?x rdf:type foaf:Person .
  ?x foaf:knows ?y .
  ?y foaf:name "Michael Douglas" .
}
```

- **Result:**

?name
"Angelina Jolie"

Presentation Outline

- The Semantic Web vision
- Linked data
- Ontologies 
- Linked geospatial and temporal data
- Applications

Ontologies

- An **ontology** is a formal, explicit, shared specification of a conceptualization of a domain (Gruber, 1993).
- **Conceptualization**: the **objects**, **concepts**, and other entities that are assumed to exist in some area of interest and the **relationships** that hold among them. A conceptualization is an abstract, simplified view of the world that we wish to represent for some purpose.
- The term **ontology** is borrowed from Philosophy, where ontology is a systematic account of existence (what things exist, how they can be differentiated from each other etc.).
- Today the word ontology is a synonym for a **shared knowledge base**.

Formal Languages for Ontologies

- Ontologies are typically expressed in some **formal logic-based language** (e.g., first-order logic).
- The literature also offers special formalisms for defining ontologies that contain mainly **taxonomic knowledge**:
 - Semantic networks
 - Frames
 - Description logics
 - RDF, RDFS and OWL in the Semantic Web.



Formal Languages for Ontologies (cont'd)

- You can think about these formalisms as being **object-oriented logics**:
 - They have special constructs for representing knowledge about **individuals** (or objects), **categories** (or classes) and **relationships** (or roles).
 - Categories are organized into **taxonomies**.
 - They have **special reasoning methods** to deal with these constructs.
- Description logics, RDF, RDFS and OWL 2 are a core part of this course.


Taxonomies

- Taxonomies have been **used profitably for centuries** in various technical fields (biology, medicine, library science etc.).
- Taxonomic information plays a central role in other **Computer Science** areas e.g., programming languages, databases and software engineering.
- Taxonomies are very important in **modern applications**: Web information retrieval and integration, knowledge management, e-commerce, e-science, e-government etc.

Example: an RDFS ontology of vehicles



Example: the GAG RDFS ontology



Very Large Ontologies and Knowledge Graphs

- Recently there has been a lot of work on developing very large ontologies and knowledge graphs that capture various areas of human knowledge and deploying this knowledge in applications such as search engines or question answering.
- Example: Watson, IBM's question answering system that beat humans in the quiz show Jeopardy (<https://www.ibm.com/history/watson-jeopardy>).

Examples of Very Large Ontologies and Knowledge Graphs

- DBpedia
- YAGO
- Freebase
- Google Knowledge Graph
- Wikidata
- NELL



- From the DBpedia Web site (<http://wiki.dbpedia.org/about>):
 - DBpedia is a community effort to extract structured information from **Wikipedia** and various Wikimedia projects so that to make this information available on the Web
 - Now, it is fed by an automated extraction process that converts structured data from Wikipedia's infoboxes into RDF triples, forming a linked data knowledge base
 - Since the emergence of the LOD Cloud in 2007, it constitutes the main resource of Linked Open Data on the Web (228+ million entities across its multilingual LOD resources)
 - The DBpedia knowledge graph has several advantages over existing knowledge graphs: it covers many domains; it represents real community agreement; it automatically evolves as Wikipedia changes, and it is truly multilingual
- SPARQL endpoint: <http://dbpedia.org/sparql>

The DBpedia Ontology

- The DBpedia ontology is the heart of Dbpedia
- Started in 2008 as a manually created ontology based on the most commonly used infoboxes within Wikipedia
- 768 classes which form a subsumption hierarchy and are described by 3000 different properties in its core schema
- 4,828,418 Resources,
- Over 900,000 Places and over 1,500,000 Persons ...



< > ↺ <https://dbpedia.org/ontology/Actor>



👁 Browse using ▾ 📄 Formats ▾

🔗 Faceted Browser

About: [actor](#)

An Entity of Type: [Class](#), from Named Graph: <http://dbpedia.org/resource/classes#>, within Data Space: dbpedia.org

An actor or actress is a person who acts in a dramatic production and who works in film, television, theatre, or radio in that capacity.

Property	Value
rdf:type	• owl:Class
rdfs:comment	• An actor or actress is a person who acts in a dramatic production and who works in film, television, the capacity. (en)
rdfs:isDefinedBy	• http://dbpedia.org/ontology/
rdfs:label	• actor (en)
rdfs:subClassOf	• dbo:Artist
owl:equivalentClass	• wikidata:Q33999
wdrs:describedby	• dbo:data/definitions.ttl
prov:wasDerivedFrom	• http://mappings.dbpedia.org/index.php/OntologyClass:Actor
is ov:defines of	• http://dbpedia.org/ontology/

DBpedia (cont'd)

- DBpedia focus on **individuals**, not general knowledge about various domains.
- DBpedia is written in RDF and it is part of the **Linked Open Data** effort.

Links to	Triples
yago	18,100,000
flickr-wrapprr	8,800,000
freebase	3,400,000
w3c-wordnet	467,101
linkedgeodata	99,075
geonames-semantic-web	86,547
dbtune-musicbrainz	22,981
opencyc	20,362
geospecies	15,972
...	...

http://dbpedia.org/page/Angelina_Jolie

File Edit View History Bookmarks Profiles Tools Help

D About: Angelina Jolie x +

dbpedia.org/page/Angelina_Jolie 70% ☆ VPN

DBpedia Browse using Formats Faceted Browser Sparql Endpoint

About: [Angelina Jolie](#)

An Entity of Type: [animal](#), from Named Graph: <http://dbpedia.org>, within Data Space: [dbpedia.org](#)

American actress



Property	Value
dbo:activeYearsStartYear	• 1982-01-01 (xsd:year)
dbo:alias	• Angelina Jolie Pitt (en)
dbo:award	• dbr:List of awards and nominations received by Angelina Jolie
dbo:birthDate	• 1975-06-04 (xsd:date)
dbo:birthName	• Angelina Jolie Voight (en)
dbo:birthPlace	• dbr:Los Angeles
dbo:birthYear	• 1975-01-01 (xsd:year)
dbo:description	• American actress (en)

10:54 AM 10/2/2026

YAGO



- From the YAGO (“Yet Another Great Ontology”) Web page (<https://yago-knowledge.org/>): YAGO is a large knowledge base with general knowledge about people, cities, countries, movies, and organizations.
- Since 2020, YAGO is built from Wikidata and the top-level taxonomy of schema.org.
- The YAGO knowledge base offers a complementary approach by prioritizing logical consistency. The KB focuses on manually defined properties, strict logical constraints, legible entity identifiers, and a curated top-level taxonomy – features that quantifiably enhance data cleanliness.
- YAGO 4.6 acts as a cleaned, logically consistent subset of Wikidata and has knowledge of more than **39 million entities** (like persons, organizations, cities, etc.)

Freebase

- From Google Research:

([https://dbpedia.org/page/Freebase_\(database\)](https://dbpedia.org/page/Freebase_(database)))

Freebase was a large collaborative knowledge base consisting of data composed mainly by its community members, including individual, user-submitted wiki contributions.

It aimed to create a global resource that allowed people (and machines) to access common information more effectively.

It was developed by the American software company Metaweb and run publicly beginning in March 2007.

Metaweb was acquired by Google.

Freebase (cont.)

- An **entity** is a single person, place, or thing. Freebase connects entities together as a **graph**.
- Freebase was bought by Google and its **Knowledge Graph** technology build on it.

Using Google's Knowledge Graph

Google search results for Leonardo da Vinci, showing the Knowledge Graph interface.

Leonardo da Vinci
Polymath

Overview | Artworks | Books | Structures | Periods | On view

Born
Apr 15, 1452
Anchiano, Italy

Died
May 2, 1519
Château du Clos Lucé, ...

Artworks

- Mona Lisa (1503)
- The Last Supper (1498)
- Salvator Mundi (1500)
- Vitruvian Man
- Virgin of the Rocks
- Lady with an Ermine (1489)

People also ask

- What is Da Vinci famous for?
- What are 5 interesting facts about Da Vinci?
- Did Leonardo da Vinci have a wife?
- What was Da Vinci's IQ?

About

Leonardo di ser Piero da Vinci was an Italian polymath of the High Renaissance who was active as a painter, draughtsman, engineer, scientist, theorist, sculptor, and architect. [Wikipedia](#)

Born: April 15, 1452, [Anchiano, Italy](#)
Died: May 2, 1519 (age 67 years), [Château du Clos Lucé, Amboise, France](#)
Periods: [Renaissance](#), [Italian Renaissance](#), [High Renaissance](#)
Place of burial: [Chapel of Saint-Hubert, Amboise, France](#)
Structures: [Leonardo's Rivellino](#)
Influenced by: [Andrea del Verrocchio](#), [Lorenzo de' Medici](#), [John Argyropoulos](#)

People also search for

- Michelangelo
- Raphael
- Leonardo DiCaprio
- Vincent van Gogh

EDO IME Wikidata

(<https://www.wikidata.org/>)

The screenshot shows the Wikidata Main Page in a web browser. The browser's address bar displays www.wikidata.org/wiki/Wikidata:Main_Page. The page features a network diagram with nodes labeled 'open', 'free', 'multilingual', 'collaborative', 'linked', and 'structured'. The central text reads 'Welcome to Wikidata' and 'the free knowledge base with 123,535,800 data entities that anyone can edit.' Below this, there are links for 'Introduction', 'Project Chat', 'Community Portal', and 'Help'. A section titled 'Want to help translate? Translate the missing messages.' is also present. The page includes a search bar, a language selector set to 'English', and a sidebar with appearance settings (Text, Width, Color). At the bottom, there are two main sections: 'Welcome!' and 'Learn about data', each with descriptive text and images. The Windows taskbar at the bottom shows various application icons and the system clock indicating 10:59 AM on 10/2/2026.

File Edit View History Bookmarks Profiles Tools Help

Wikidata

www.wikidata.org/wiki/Wikidata:Main_Page

WIKIDATA Search Wikidata Search

English Donate Create account Log in

Main Page Discussion Read View source View history

collaborative linked

open

free multilingual

structured

Welcome to Wikidata

the free knowledge base with 123,535,800 data entities that anyone can edit.

Introduction • Project Chat • Community Portal • Help

Want to help translate? [Translate the missing messages.](#)

Welcome!

Wikidata is a free and open knowledge base that can be read and edited by both humans and machines.

Wikidata acts as central storage for the **structured data** of its Wikimedia sister projects including Wikipedia, Wikivoyage, Wiktionary, Wikisource, and others.

Wikidata also provides support to many other sites and services beyond just Wikimedia projects! The content of Wikidata is [available under a free license](#), exported using [standard formats](#), and can be [interlinked to other open data sets](#) on the linked data web.

Learn about data

New to the wonderful world of data? [Develop and improve your data literacy through content](#) designed to get you up to speed and feeling comfortable with the fundamentals in no time.

10:59 AM 10/2/2026

Wikidata



- Wikidata is Wikimedia project like Wikipedia.
- It is a free and open knowledge base constructed mostly manually through a community process.
- In December of 2014, Google closed down Freebase and offered its content to Wikidata.
- See the following paper for a discussion of the migration process:
<https://static.googleusercontent.com/media/research.google.com/en//pubs/archive/44818.pdf>

NELL (Never-Ending Language Learning)

Read the Web

Research Project at Carnegie Mellon University

- Home
- Project Overview
- Resources & Data
- Publications
- People

NELL: Never-Ending Language Learning

Can computers learn to read? We think so. "Read the Web" is a research project that attempts to create a computer system that learns over time to read the web. Since January 2010, our computer system called NELL (Never-Ending Language Learner) has been running continuously, attempting to perform two tasks each day:

- First, it attempts to "read," or extract facts from text found in hundreds of millions of web pages (e.g., `playsInstrument(George_Harrison, guitar)`).
- Second, it attempts to improve its reading competence, so that tomorrow it can extract more facts from the web, more accurately.

So far, NELL has accumulated over 50 million candidate beliefs by reading the web, and it is considering these at different levels of confidence. NELL has high confidence in 4,010,713 of these beliefs — these are displayed on this website. It is not perfect, but NELL is learning. You can track NELL's progress below or [@cmunell on Twitter](#), browse and download its [knowledge base](#), read more about our [technical approach](#), or join the [discussion group](#).

Recently-Learned Facts

Refresh

instance	iteration	date learned	confidence	
washtenaw_lotus is an automobile manufacturer	1075	04-sep-2017	98.9	 
estate_resolutions is a winery	1075	04-sep-2017	96.3	 

Interacting with NELL

NELL Knowledge Base Browser
CMU Read the Web Project

log in | preferences | help/instructions | feedback

categories

relations

- everypromotedthing
- abstractthing
 - event
 - convention
 - musicfestival
 - protestevent
 - meetingeventtitle
 - conference
 - mlconference
 - weatherphenomenon
 - sportsevent
 - sportsgame
 - race
 - olympics
 - grandprix
 - crimeorcharge
 - earthquakeevent
 - election
 - bombingevent
 - militaryeventtype
 - militaryconflict
 - productlaunchevent
 - filmfestival
 - roadaccidentevent
 - meetingeventtype
 - eventoutcome
 - mlalgorithm
 - physiologicalcondition
 - disease
 - nondiseasecondition
 - religion
 - creativework
 - musicalbum
 - book
 - poem
 - lyrics
 - visualartform
 - movie

perl_php (programminglanguage)
literal strings: [perl_php](#)

Help NELL Learn!

NELL wants to know if this belief is correct.
If it is or ever was, click thumbs-up. Otherwise, click thumbs-down.

- [perl_php](#) is a [programming language](#) 👍👎

categories

- [programminglanguage](#)(100.0%)
 - LE @1075 (100.0%) on 02-sep-2017
 - CPL @1058 (71.8%) on 16-may-2017 ["languages such as _" "development languages such as _"] using perl_php

Building Ontologies

- The example ontologies and knowledge graphs we saw have been created using different means:
 - Getting **a team of domain experts and ontology engineers** to organize the ontology and write the axioms (WordNet).
 - **Importing structured facts** from various Web sources possibly with some inference (DBpedia, FreeBase, YAGO2). Manual curation is also possible here (FreeBase, Wikidata).
 - Parsing **text (web) documents** and extracting information from them.

- The availability of open sources of knowledge on the Web has fuelled the development of open source and community efforts such as DBpedia, FreeBase and YAGO. **We expect this trend to continue in the future!**
- In addition industrial strength commercial knowledge graphs have been developed by big Web companies such as Google, Microsoft, Facebook etc.
- **Ontology engineering** offers us general methodologies of how to create, maintain and evolve ontologies.

Description Logics

- The origins of DLs lie in research on semantic networks and frames. DLs are languages for describing the **nature and structure of objects**.
- The DL approach to KR was developed in the 80's and 90's in parallel with pure FOL approaches and other languages for structured objects like Telos and F-logic.
- Recently, DLs have been used to provide the **foundations for ontology languages for the Web** e.g., OWL.
- DLs are logics based on **descriptions of concepts or terms**.

Examples

- The set of female doctors.

Female $\sqcap$ *Doctor*

- The set of individuals that have at least 3 children that are male.

(≥ 3 hasChild . Male)

- The set of individuals such that all their children have graduated from at least one Greek University.

($\forall$ hasChild . ($\exists$ isAlumniOf . GreekUniversity))

- The set of individuals that have at least three children such that all their degrees are from Greek Universities.

$(\geq 3 \text{ hasChild} . (\forall \text{isAlumniOf} . \text{GreekUniversity}))$

- Ann is a female doctor.

(Female $\sqcap$ Doctor)(ANN)

- John is a child of Ann.

hasChild(ANN,JOHN)

- Ann has at least 3 children that are male.

(≥ 3 hasChild . Male)(ANN)

- A woman is a female person.

$$Woman \equiv Person \sqcap Female$$

- A parent is a person who has at least one child.


$$Parent \equiv Person \sqcap (\exists hasChild . Person)$$

OWL

- OWL 2 is the current version of the **Web Ontology Language** and a W3C recommendation as of October 2009.
- The previous version of OWL (OWL 1) became a W3C recommendation in 2004.
- All W3C documents about OWL 2 can be found at <https://www.w3.org/TR/owl2-overview/>.



The Structure of OWL 2



Example

```
SubClassOf (:Student :Person)
```

```
EquivalentClasses (:Man  
    ObjectIntersectionOf (:Person :Male))
```

```
EquivalentClasses (:Woman  
    ObjectIntersectionOf (:Person :Female))
```

```
EquivalentClasses (:Parent  
    ObjectSomeValuesFrom (:hasChild :Person))
```

Example (cont'd)

```
EquivalentClasses (:Teenager  
  ObjectIntersectionOf (Person  
    DataSomeValuesFrom (:hasAge  
      DatatypeRestriction (xsd:integer  
        xsd:minExclusive "12"^^xsd:integer  
        xsd:maxInclusive "19"^^xsd:integer ) ) ) )
```

Example (cont'd)

```
EquivalentClasses (:ChildlessPerson  
  ObjectIntersectionOf (:Person  
    ObjectComplementOf (ObjectSomeValuesFrom (  
      ObjectInverseOf (:hasParent) owl:Thing)))
```

Example (cont'd)

```
ClassAssertion(:Person :John)
```

```
ClassAssertion(:Male :John)
```

```
ClassAssertion(:Person :Mary)
```

```
ClassAssertion(:Female :Mary)
```

```
ObjectPropertyAssertion(:hasChild :John :Mary)
```

```
NegativeObjectPropertyAssertion(:hasChild :John :Tina)
```

```
DataPropertyAssertion(:hasAge :Maria "18"^^xsd:integer)
```

Example (cont'd)


SameIndividual (:John :Jack)

SameIndividual (:John otherOnt:JohnBrown)

SameIndividual (:Mary otherOnt:MaryBrown)

DifferentIndividuals (:John :Bill)


Presentation Outline

- The Semantic Web vision
- Linked data
- Ontologies
- Linked geospatial and temporal data 
- Applications


Geospatial Data on the Web



Google™ earth



Earth Observation Data



Linked Geospatial and Temporal Data

- This is a research area studied by prof. Koubarakis' group since 2010.
- The basic research question is **how to manage geospatial and temporal data on the Web using linked data technologies.**
- <http://ai.di.uoa.gr/>


Research Contributions

- The model stRDF and stSPARQL
- The systems GeoTriples and Silk
- The spatiotemporal RDF store Strabon
- The ontology-based data access system Ontop-spatial.
- The visualization tool Sextant.
- Many applications, especially with Earth Observation data.
- See the survey paper:
<http://cgi.di.uoa.gr/~koubarak/publications/2017/ic4Id-manolis.pdf>
- See also:
<http://cgi.di.uoa.gr/~koubarak/publications/2018/template-based-GeoQA.pdf>

Example of stRDF (Geospatial Dimension)

```
gag:Olympia
  rdf:type gag:MunicipalCommunity;
  gag:name "Ancient Olympia";
  gag:population "184"^^xsd:int;
  strdf:hasGeometry "POLYGON((21.5 18.5,23.5 18.5, 23.5
21,21.5 21,21.5 18.5));
  http://www.opengis.net/def/crs/OGC/1.3/CRS84"^^strdf:WKT.
```

Ancient Olympia



Example of stSPARQL (Geospatial Dimension)

Query: Compute the parts of burnt areas that lie in coniferous forests

```
SELECT ?burntArea (strdf:intersection(?baGeom,  
                                strdf:union(?fGeom)) AS ?burntForest)  
WHERE {  ?burntArea  rdf:type  noa:BurntArea;  
                                strdf:hasGeometry ?baGeom.  
          ?forest  rdf:type  clc:Region;  
                                clc:hasLandCover  clc:ConiferousForest;  
                                strdf:hasGeometry ?fGeom.  
  
FILTER    (strdf:intersects(?baGeom,?fGeom)) }  
GROUP BY ?burntArea ?baGeom
```



Presentation Outline

- The Semantic Web vision
- Linked data
- Ontologies
- Linked geospatial and temporal data
- Applications 

Applications

- Data integration (e.g., see project Optique <http://www.optique-project.eu/>)
- E-government (e.g., open data)
- E-commerce
- Tourism
- Medicine
- Biology
- Earth Observation (see the work of Prof Koubarakis' group in projects TELEIOS <http://www.earthobservatory.eu/> and LEO <http://www.linkedeodata.eu/>).
- ...

Readings on the SW and Linked Data

- The papers on the Semantic Web vision mentioned in the presentation.
- Chapter 1 of the Semantic Web Primer by Grigoris Antoniou and Frank van Harmelen available from <http://www.csd.uoc.gr/~hy566/SWbook.pdf>.
- Chapter 1 of the book “Foundations of Semantic Web Technologies” by Pascal Hitzler, Markus Krötzsch and Sebastian Rudolph. See <http://semantic-web-book.org>.
- The book “Linked Data: Evolving the Web into a Global Data Space” by Tom Heath and Christian Bizer available from <http://linkeddatabook.com/editions/1.0/>.
- The original W3C Semantic Web Activity web site (<http://www.w3.org/2001/sw/>). Browse! There are many interesting introductory tutorials. (New site: <https://www.w3.org/2013/data/>)
- The ESWC invited talk of Jim Hendler available at http://videolectures.net/eswc2011_hendler_work/ which looks at what the Semantic Web has achieved so far and deals with various criticisms of the area.