



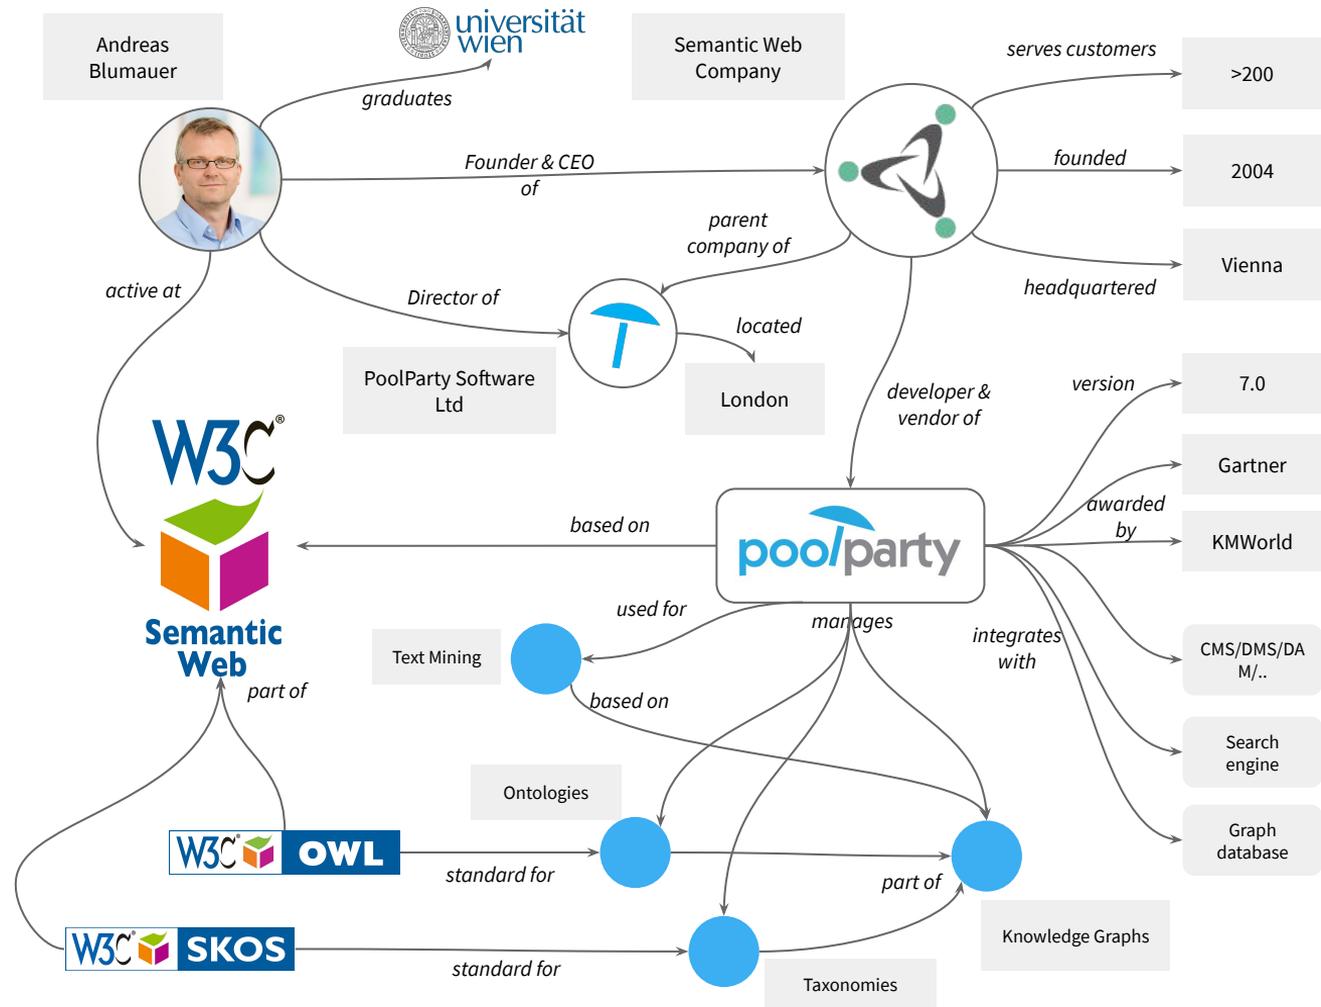
Introduction to Knowledge Graphs

When data learn to think outside the box

Andreas Blumauer
CEO of Semantic Web Company



Introduction



WHAT'S THE PROBLEM?

Many Challenges in (Enterprise) Information Management & AI

“Search” is still about documents only

... and in enterprises it's painful

wind farms OR wind parks OR wind power plants OR wind power station: 

Different shades of metadata



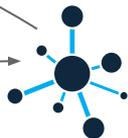
User-agnostic and context-free data models

Do machines understand user intent? Do they have enough context?

1. Intent recognition
2. Entity linking
3. Background knowledge

French SUV

Which Sport-utility vehicle from France provides enough space for my family with 3 kids?

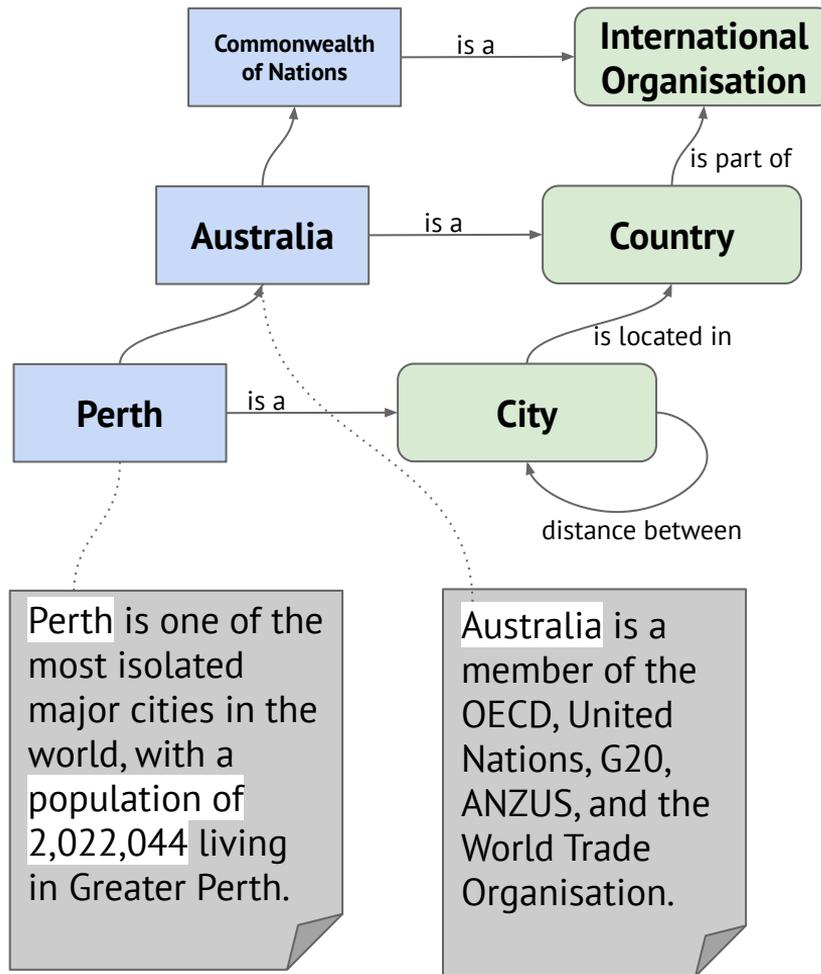


The Peugeot 5008 breaks new ground as a large SUV with many features.

Car	Loadspace	Max no. of seats
KIA Sorento	550 litres	7
Peugeot 5008	823 litres	7
BMW X3	550 litres	5

Artificial “Intelligence”

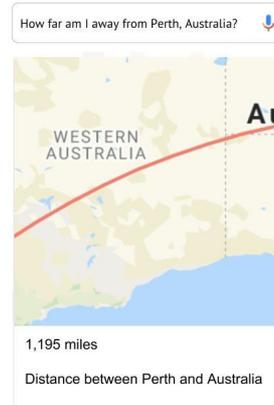
Background knowledge is key



Support complex Q&A:

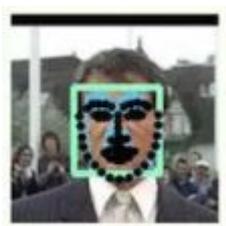
Which cities located in the Commonwealth of Nations have a population of more than 2 mio. people?

Avoid illogical answers:



Machine Learning per Data Silo

Lack of AI Strategy



Deep Learning
Genetic Algorithms



Monte Carlo TS
Deep Learning



Neuronal
networks
Case based
reasoning

Who is **CEO** of a **Bank**,
headquartered in Europe that
generates revenue per employee
higher than 400,000 Euro?



Desktop Data Integration

Knowledge workers spend a large part of their working time on ad hoc data integration tasks, the so-called "research".

Which parts of it could be automated?



	Head-quarter	Ticker Symbol	Revenue
Credit Suisse	Zurich	VTX: CSGN	CHF 23.4b
HSBC	London	LON: HSBA	USD 60.0b
Allianz	Munich	ETR: ALV	EUR 122.3b
Deutsche Bank	Frankfurt	ETR: DBK	EUR 33.5b



```
<Employees>
<VTX:CSGN>
  48,200
</VTX:CSGN>
<LON:HSBA>
  266,273
</LON:HSBA>
<ETR:ALV>
  147,425
</ETR:ALV>
<ETR:DBK>
  101,104
</ETR:DBK>
</Employees>
```



The task awaiting Tidjane Thiam when he takes over from Brady Dougan as the new chief executive at Credit Suisse Group AG is clear: how to pull the Swiss bank out of a post-financial crisis rut.



Many challenges in data management

- ▶ Heterogeneous data
- ▶ Implicit semantics
- ▶ No background knowledge
- ▶ Data quality difficult to measure
- ▶ Proprietary schemas
- ▶ Unstructured data
- ▶ Ambiguity
- ▶ Multilinguality
- ▶ Various units and currencies

→ From Search to QA engines

	Head-quarter	Ticker Symbol	Umsatz
Credit Suisse	Zürich	VTX: CSGN	CHF 23.4b
HSBC	London	LON: HSBA	USD 60.0b
Allianz	München	ETR: ALV	EUR 122.3b
Deutsche Bank	Frankfurt	ETR: DBK	EUR 33.5b



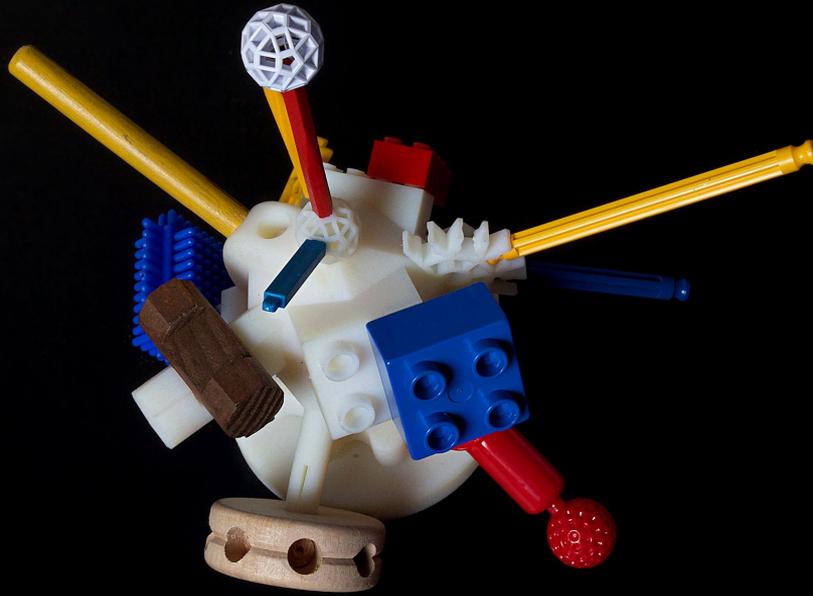
```

<Employees>
  <VTX:CSGN>
    48,200
  </VTX:CSGN>
  <LON:HSBA>
    266,273
  </LON:HSBA>
  <ETR: ALV>
    147,425
  </ETR: ALV>
  <ETR: DBK>
    101,104
  </ETR: DBK>
</Employees>

```

The task awaiting Tidjane Thiam when he takes over from Brady Dougan as the new chief executive at Credit Suisse Group AG is clear: how to pull the Swiss bank out of a post-financial crisis rut.

The Free Universal Construction Kit is a matrix of adapter bricks that enable complete interoperability between ten popular children's construction toys.



<http://www.free-universal-construction-kit.com>
By Golan Levin and Shariq Sanyal for the Free Art & Technology (FAAT) Lab and SynLab.net, 2012

The FREE UNIVERSAL CONSTRUCTION KIT is a matrix of adapter bricks that enable complete interoperability between ten popular children's construction toys. By allowing any piece to join to any other, the Kit encourages totally new forms of intercourse between otherwise closed systems. As with other grassroots interoperability remedies, the Kit implements proprietary protocols in order to provide a public service unmet, or unmeetable, by corporate interests.



Free Universal Construction Kit (Golan Levin et al)

‘Interoperability’ begins with people

Bringing various types of mindsets together

Data-centric people	Knowledge-centric people
Service-oriented industries	Asset-oriented industries
Databases & Excel	Content & documents
Structured data	Unstructured data
Linked Data	Knowledge Graph
Algorithms	Collaborative knowledge management
Predictability and automation	Product innovation and risk mitigation
Ontologies and Machine Learning	Taxonomies and Collaboration
“Actionable Data!”	“Better Decisions!”
Understand the customer and market place better	Stay or become the expert in the field

What is a Knowledge Graph?

From a Knowledge Engineer's perspective

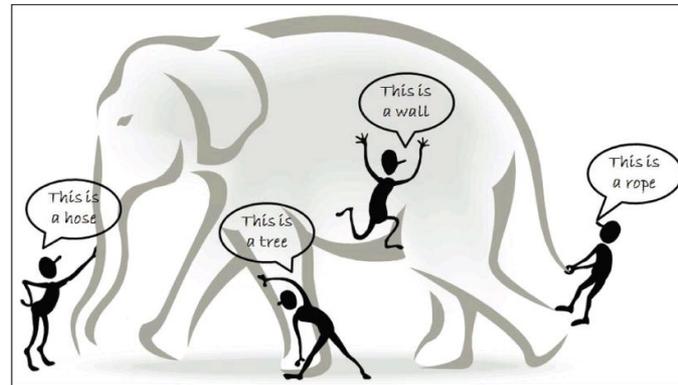
A Knowledge Graph is a **model of a knowledge domain** created by subject-matter experts with the help of intelligent machine learning algorithms.

From a Data Architect's perspective

Structured as an additional virtual data layer, the KG lies on top of existing databases or data sets to **link all your data together at scale** – be it structured or unstructured.

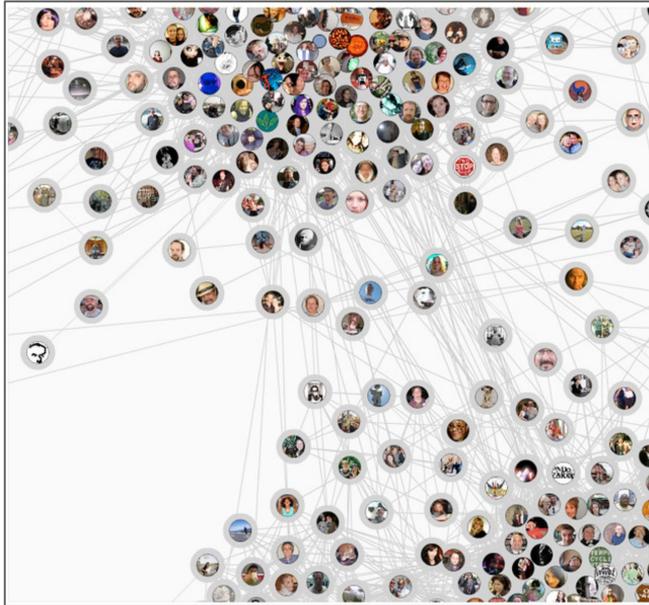
From a Data Engineer's perspective

It provides a **structure and common interface** for all of your data and enables the creation of smart multilateral relations throughout your databases.

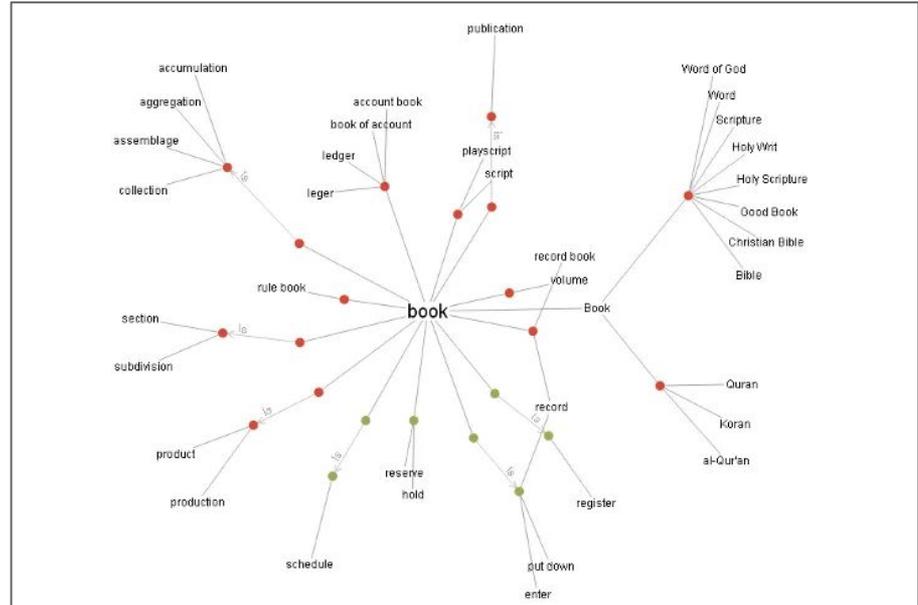


Other Graphs than *Knowledge* Graphs

Mind the gap!



Social Graphs / Networks



Wordnets



Stardog



Oracle Spatial & Graph

- ▶ Amazon Neptune
- ▶ Marklogic
- ▶ AllegroGraph
- ▶ GraphDB
- ▶ Azure Cosmos DB
- ▶ Neo4j
- ▶ ...

Which Graph Model?

RDF versus
Property Graphs

→ Converging approaches

	Property Graph	RDF Graph (Triple Stores)
<i>Main use case</i>	Traverse a graph	Query a graph
<i>Typical applications</i>	Path Analytics, Social Network Analysis	Data Integration, Knowledge Representation
<i>Standards</i>	No standards → Gremlin, Cypher, PGQL, ...	W3C Semantic Web standards → SPARQL 1.1
<i>Additional options (Example)</i>	Shortest path calculations	Inferencing

Sounds interesting? Gimme some examples!

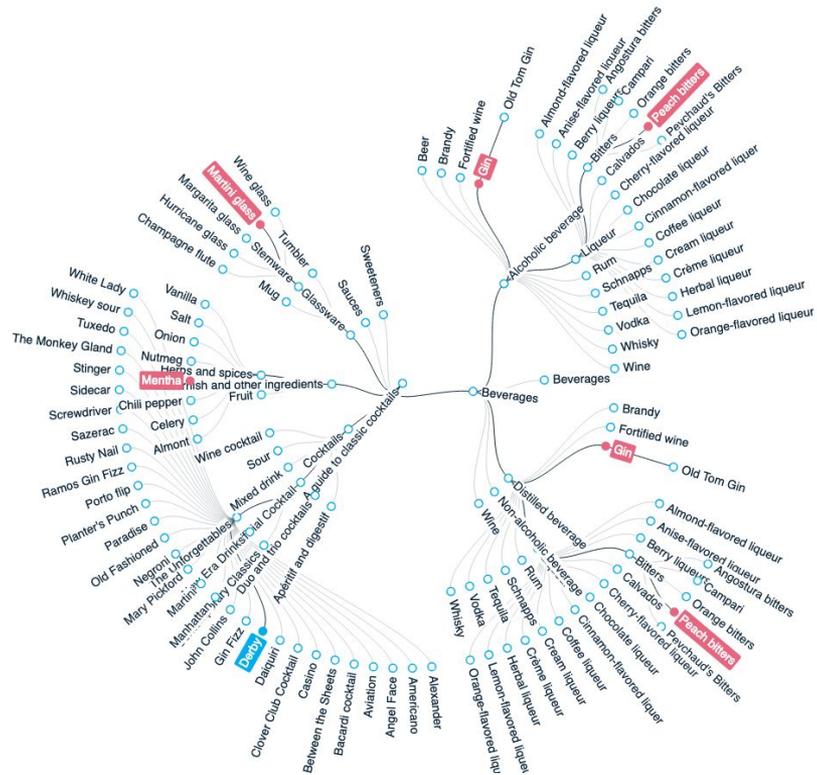
WHO USES KGs?

An overview over typical KG use cases

Can you see the Knowledge Graph?



[Taobao \(Alibaba\)](#)



[PoolParty GraphViews](#)

Google Knowledge Graph

About 87,000,000 results (0.46 seconds)

PoolParty Semantic Suite - Your Complete Semantic Platform

<https://www.poolparty.biz>

PoolParty Semantic Suite uses innovative means to help organizations build and manage enterprise knowledge graphs as a basis for various AI applications.

[Product Overview Page](#)

PoolParty Semantic Suite is a semantic platform that helps ...

[Product Matrix & Prices](#)

Get an overview over different PoolParty Semantic Suite ...

[PoolParty Thesaurus Server](#)

PoolParty Thesaurus Server. Enterprise Taxonomy and ...

[More results from poolparty.biz](#)

[PoolParty Semantic Integrator](#)

Agile Data Integration based on PoolParty Semantic Integrator ...

[Vocabulary Hub](#)

Vocabulary Hub containing publicly available SKOS ...

[PoolParty PowerTagging](#)

PowerTagging for SharePoint 2013, SharePoint 2016 and ...

PoolParty Semantic Suite - Wikipedia

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

The PoolParty Semantic Suite is a technology platform provided by the Semantic Web Company. The EU-based company belongs to the early pioneers of the ...

[History](#) · [Product](#) · [Technologies](#) · [Awards & Recognition](#)

PoolParty Semantic Suite - Semantic Web Company

<https://semantic-web.com/poolparty-semantic-suite>

PoolParty semantic suite is world-class technology for your data. It offers sharply focused solutions for knowledge organization and content business.

PoolParty Semantic Suite

Software



The PoolParty Semantic Suite is a technology platform provided by the Semantic Web Company. The EU-based company belongs to the early pioneers of the Semantic Web movement. The software supports enterprises in knowledge management, data analytics and content organisation. [Wikipedia](#)

Stable release: 7

Operating system: [Windows](#), [Mac](#), [Linux](#)

Developer(s): [Semantic Web Company](#)

Written in: [Java](#)

[Feedback](#)

Knowledge Graph Adoption

Largest change in market cap by company (2009 to 31 March 2018)

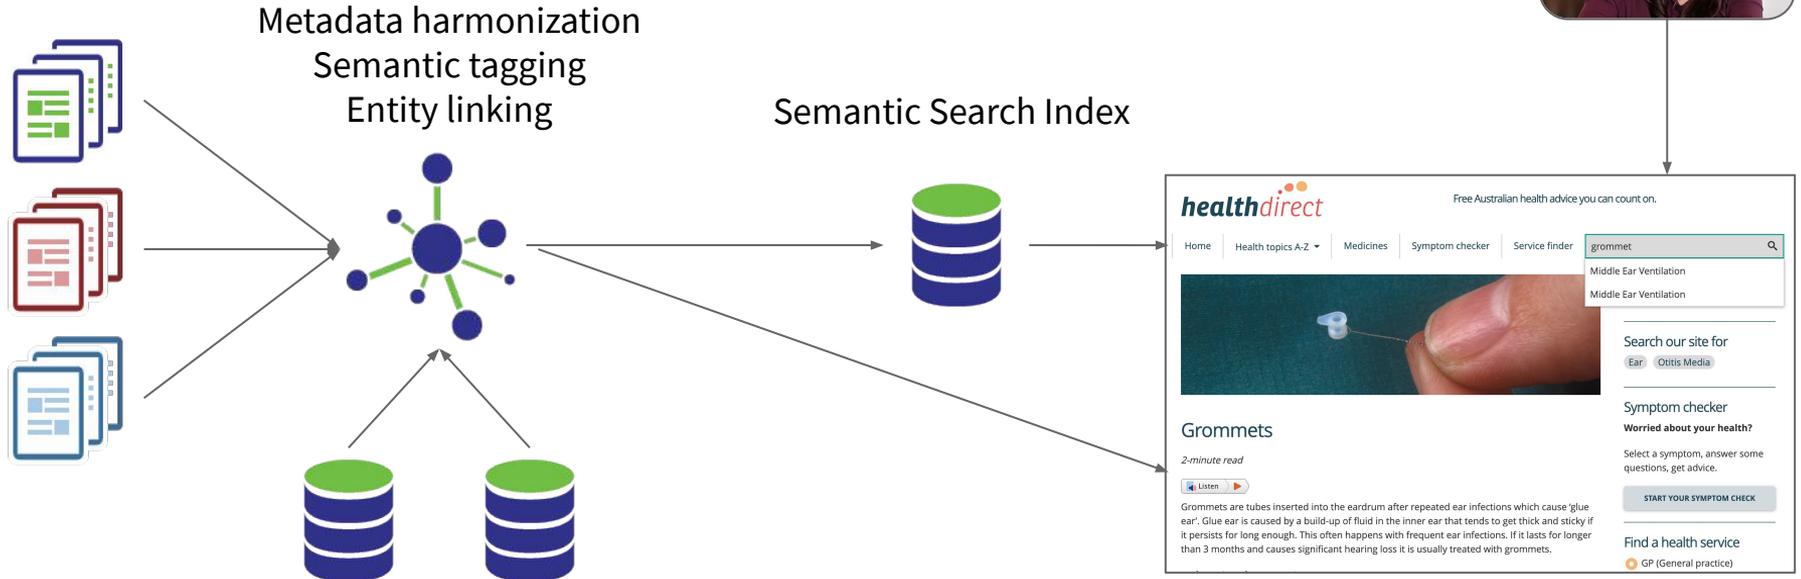
	Company name	Location	Industry	Change in market cap 2009-2018 (\$bn)	Market cap 2018 (\$bn)
Known knowledge graph builders	1 Apple	United States	Technology	757	851
	2 Amazon.Com	United States	Consumer Services	670	701
	3 Alphabet	United States	Technology	609	719
	4 Microsoft Corp	United States	Technology	540	703
	5 Tencent Holdings	China	Technology	483	496
	6 Facebook	United States	Technology	383(1)	464
Operator of Taobao and KG builder	7 Berkshire Hathaway	United States	Financial	358	492
Known KG builders	8 Alibaba	China	Consumer Services	302(1)	470
	9 JPMorgan Chase	United States	Financials	275	375
	10 Bank of America	United States	Financials	263	307

Source: [Collapsing the IT Stack: Clearing a path for AI adoption](#)
Alan Morrison (Sr. Research Fellow at PwC)

Knowledge Graphs for **Information Retrieval**

Semantic tagging, query expansion, faceted search, classification, similarity-based recommender

Now I can find all the documents in one place, and get assistance with it.



Example healthdirect Australia

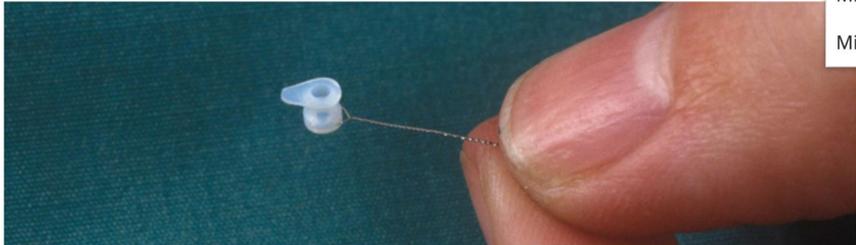

Free Australian health advice you can count on.

Home
Health topics A-Z ▾
Medicines
Symptom checker
Service finder

🔍

Middle Ear Ventilation

Middle Ear Ventilation



Grommets

2-minute read

🔊 Listen
▶

Grommets are tubes inserted into the eardrum after repeated ear infections which cause 'glue ear'. Glue ear is caused by a build-up of fluid in the inner ear that tends to get thick and sticky if it persists for long enough. This often happens with frequent ear infections. If it lasts for longer than 3 months and causes significant hearing loss it is usually treated with grommets.

Search our site for

Ear
Otitis Media

Symptom checker

Worried about your health?

Select a symptom, answer some questions, get advice.

START YOUR SYMPTOM CHECK

Find a health service

○
GP (General practice)

<https://www.healthdirect.gov.au/>

Semantic Search

Renewable Energy Glossary

- energy (12)
- wind (7)
- windpower (6)
- wind farms (5)
- wind turbines (5)
- generators (4)
- photovoltaic power (4)
- ocean energy (2)

More

Climate Compatible Development Glossary

- renewable energies (6)
- wind power (6)
- projects (4)
- economic cost (3)
- fossil energy (3)
- greenhouse gas emissions (3)

wind farms

Show tagged content only



Wind farms

[SHOW PARENT](#)

A wind farm is a group of wind turbines in the same location used for production of ... For example, Gansu Wind Farm, the largest wind farm in the world, has several thousand turbines ...

WIND FARMS



Basics of Wind Energy

[SHOW PARENT](#)

What is a wind farm ... The turbines in a wind farm are connected so the electricity they generate can travel from the wind farm to the power grid ...

ELECTRICITY GENERATION ENERGY GRIDS WIND WIND POWER WIND TURBINES WINDPOWER



Wind_test

[SHOW PARENT](#)

Wind farms consist of many individual wind turbines which are connected to the electric ... Small onshore wind farms can feed some energy into the grid or provide electricity to ...

ECONOMIC COST ELECTRICITY GENERATION ENERGY WIND WIND FARMS WIND POWER WIND TURBINES

WINDPOWER



France Gears Up for Floating Wind

[SHOW PARENT](#)

This project intends to quickly put floating wind farms on the map as a competitive energy ... Engie has stated: "Floating wind turbines are an up-and-coming technology that can be ...

wind farms



Description

A group of wind turbines interconnected to a common power provider system through a system of transformers, distribution lines, and (usually) one substation. Operation, control, and maintenance functions are often centralized through a network of computerized monitoring systems, supplemented by visual inspection.

Also known as

wind parks, wind power plants, wind power stations, windfarm

More general

[windpower concepts](#)

More specific

[grid-connected wind power systems, on-shore windparks, off-shore windparks](#)

Related

[wind turbines, wind power capacity installed, wind feed-in tariffs, stand-alone wind power turbines](#)

PowerTagging and PowerSearch for SharePoint

Use Case Research in Life Sciences

As a researcher in pharmaceutical industry, I want to plan new experiments more efficiently. I want to know what's already available. I'm interested in former experiments where

- certain genes were tested
- under specific treatment conditions
- in a target therapeutic area
- with help from categorisation systems like 'disease hierarchies'

→ Linking Structured to Unstructured Data and to Industry Knowledge Graphs

UniProt, ChEMBL



Experiments
Documentation



MeSH

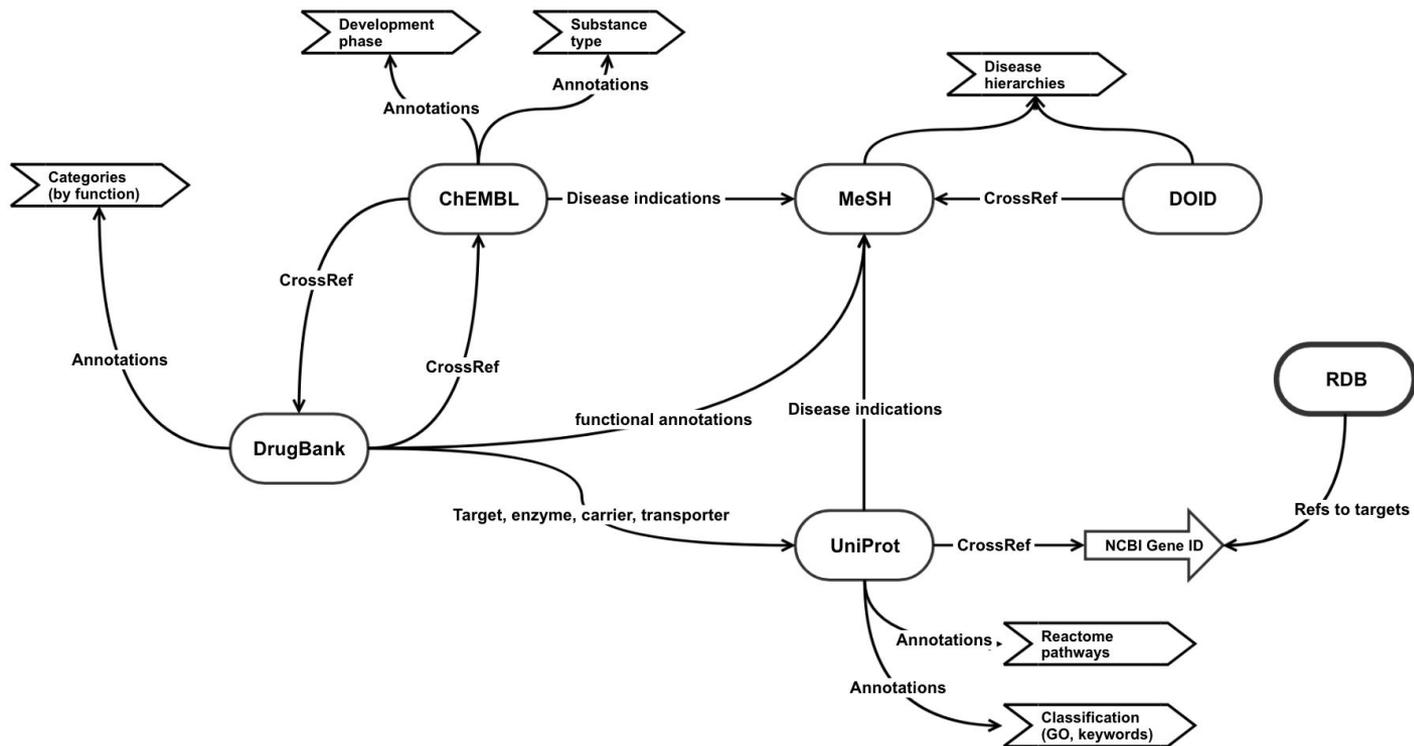


DrugBank



Reusing Graphs

Knowledge Graphs used as a Semantic Search Index

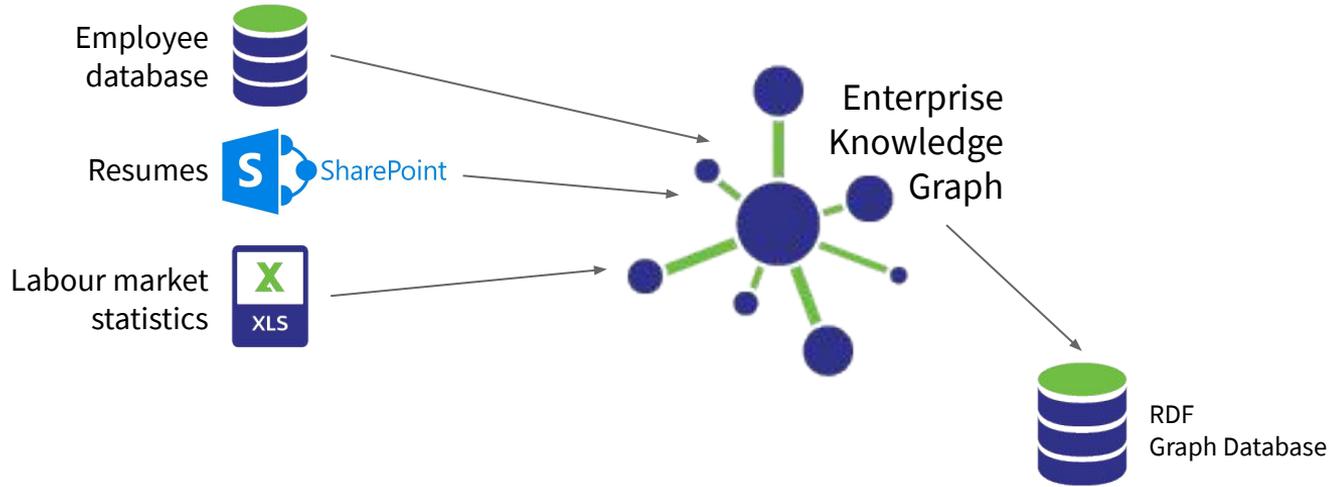


Gartner - Hype Cycle for AI (2018):

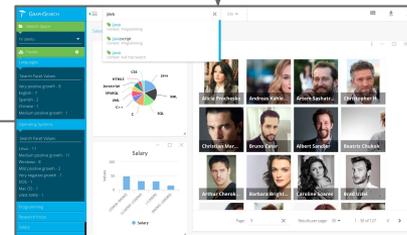
“Knowledge Graphs serve as means to enrich unstructured information to provide a rich set of additional access points to document repositories”

Knowledge Graphs for **Data Integration & Analytics**

Metadata enrichment, linked data, text mining, entity-centric search, agile reporting

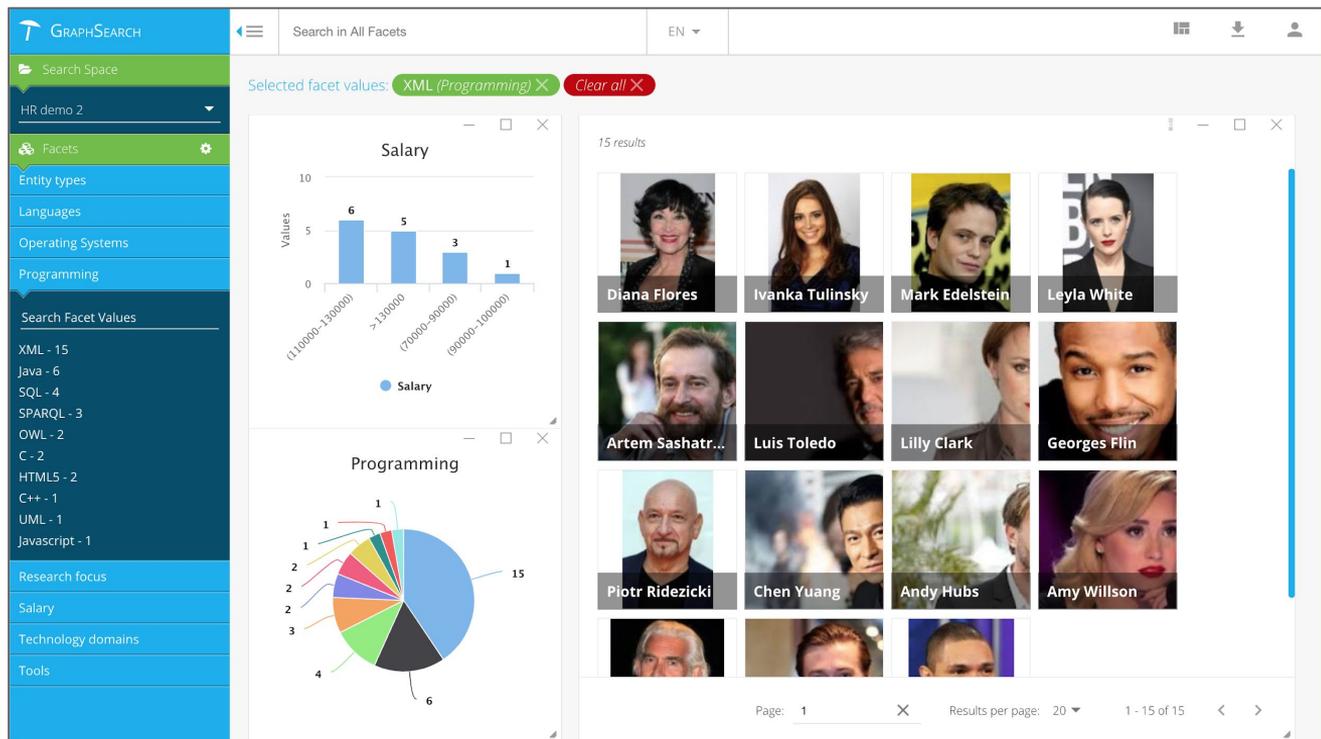


Now I can identify employees along many dimensions.



PoolParty GraphSearch

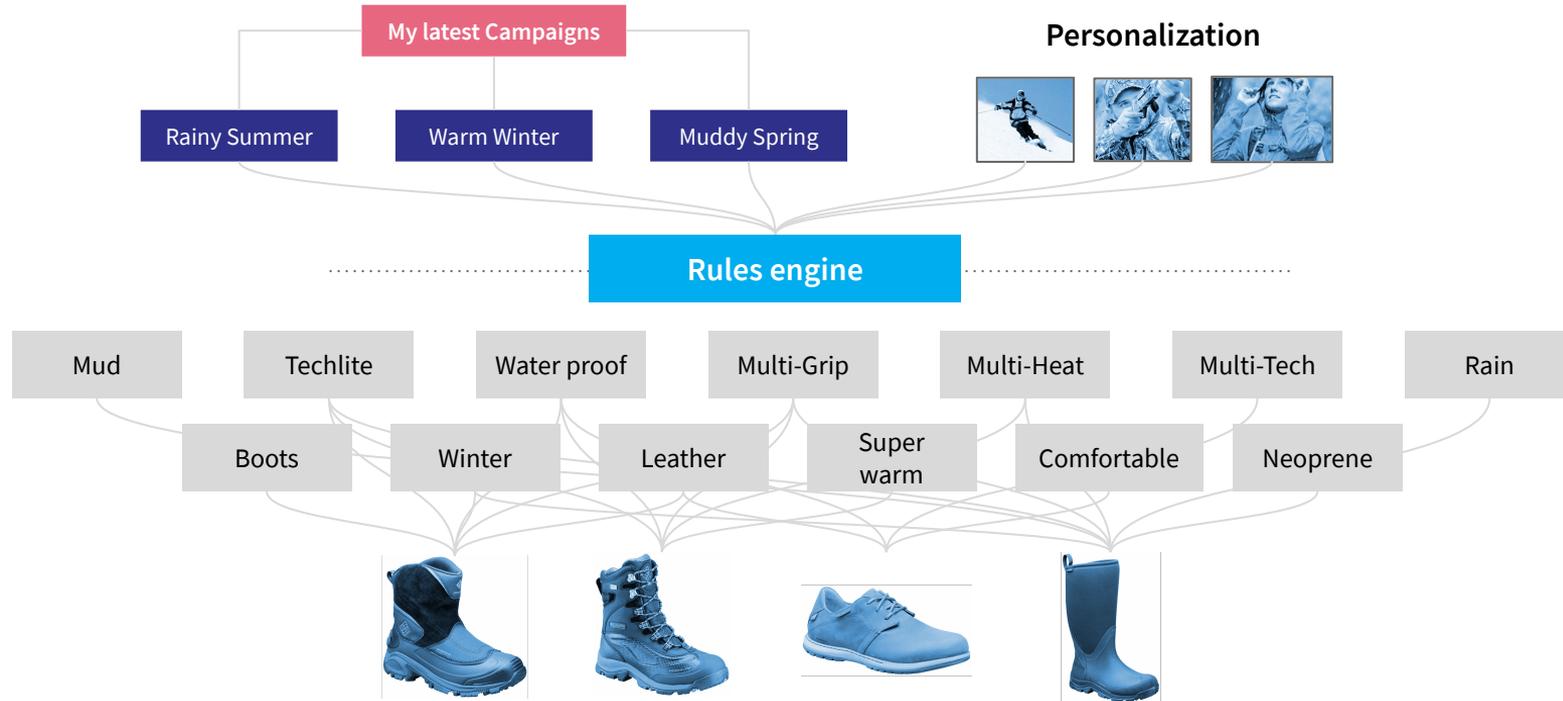
Entity-centric Search & Analytics



HR Analytics Dashboard

Knowledge Graphs for E-commerce/Content Mgmt.

Personalization, recommender / Dynamic menus, campaigns / Translating jargon into client's needs



Product Recommender

Parrano, UnieKaas
✕

←
→
Bannert Cuvee Schatzberg 2...
> Parrano, UnieKaas



Parrano is aged for five months, developing its nutty Parmigiano-Reggiano flavors while maintaining the firm, smooth texture of a young Gouda. This perfect combination brings you the best of both worlds: a cheese that is full of flavor and versatile enough to cut, grate and melt

CHEESE ▼

CHEESE DESCRIPTORS ▲

- Creamy
- Higher fat
- Semi-hard cheese
- Salty

See Also >

Wine Cheese Similarity ▼


Parrano ch...


Gaskaas P...


Remeker P...


Le Petit Do...


dse Komijn ...


Maasdam...





Similarities: 10 ▼

Recommended Wine Cheese Recommender ▼


Domäne W...


Tauss Blau...


Prieler Leit...


Bannert C...


Uwe Schief...


Moric - Bla...


Pinot gris


Weninger ...


Allram Gr...

Recommendations: 10 ▼

CLOSE

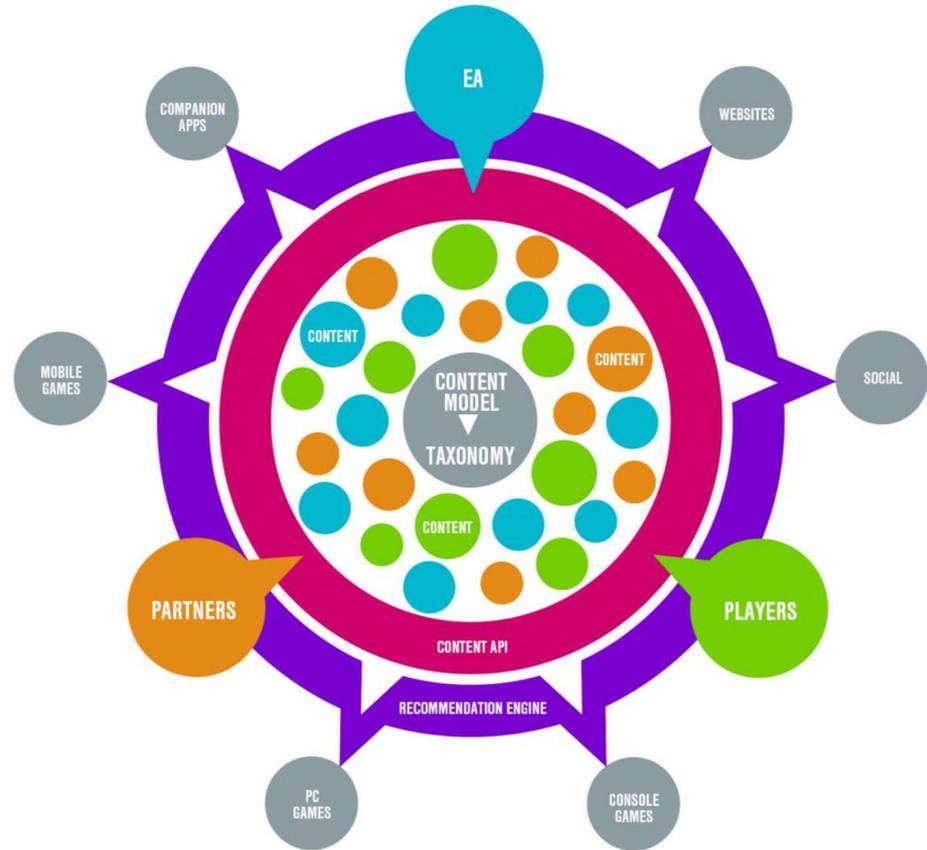
Wine and Cheese Harmonizer

© Semantic Web Company 2019

28

Example

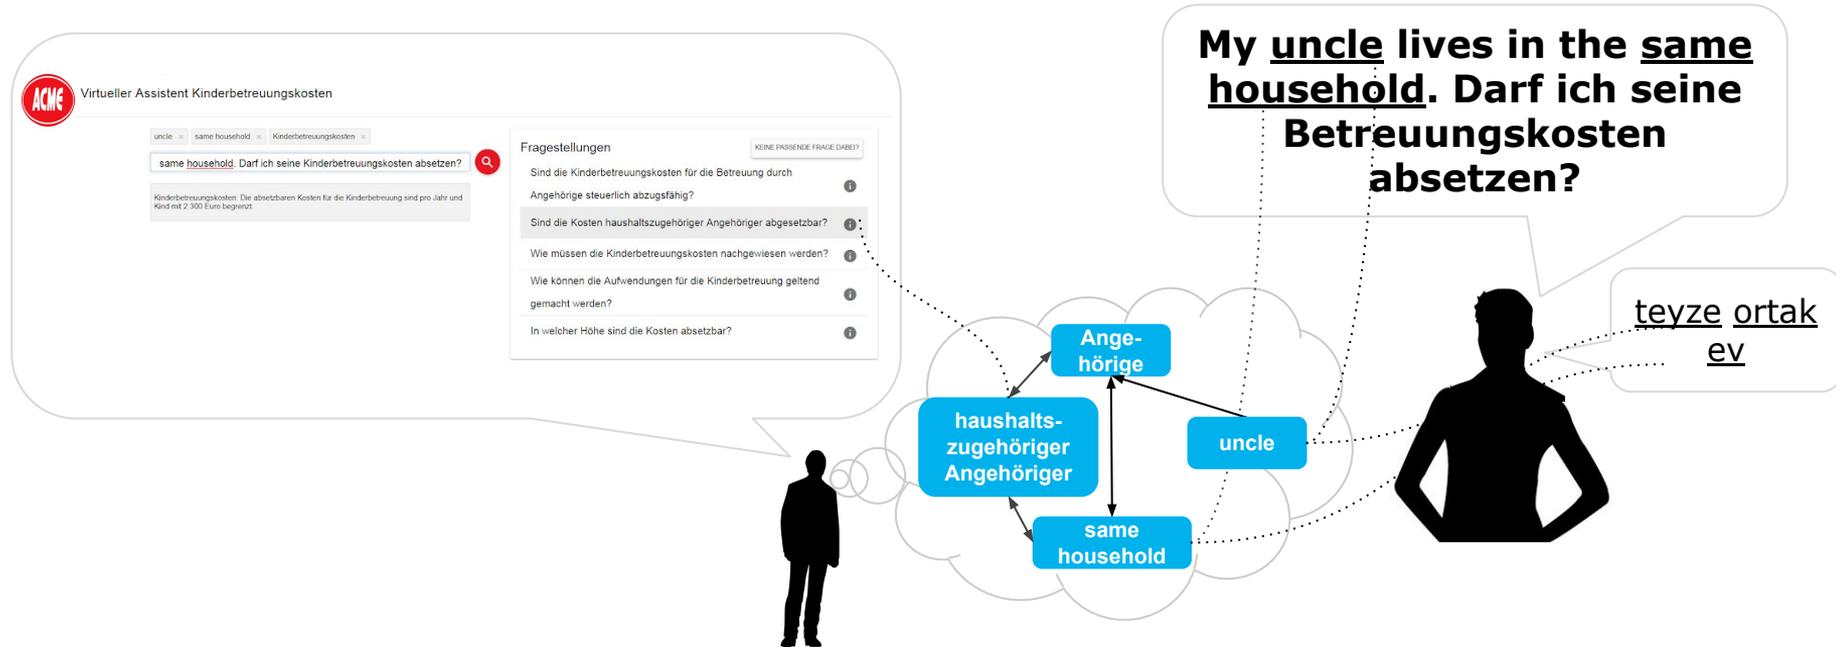
Electronic Arts



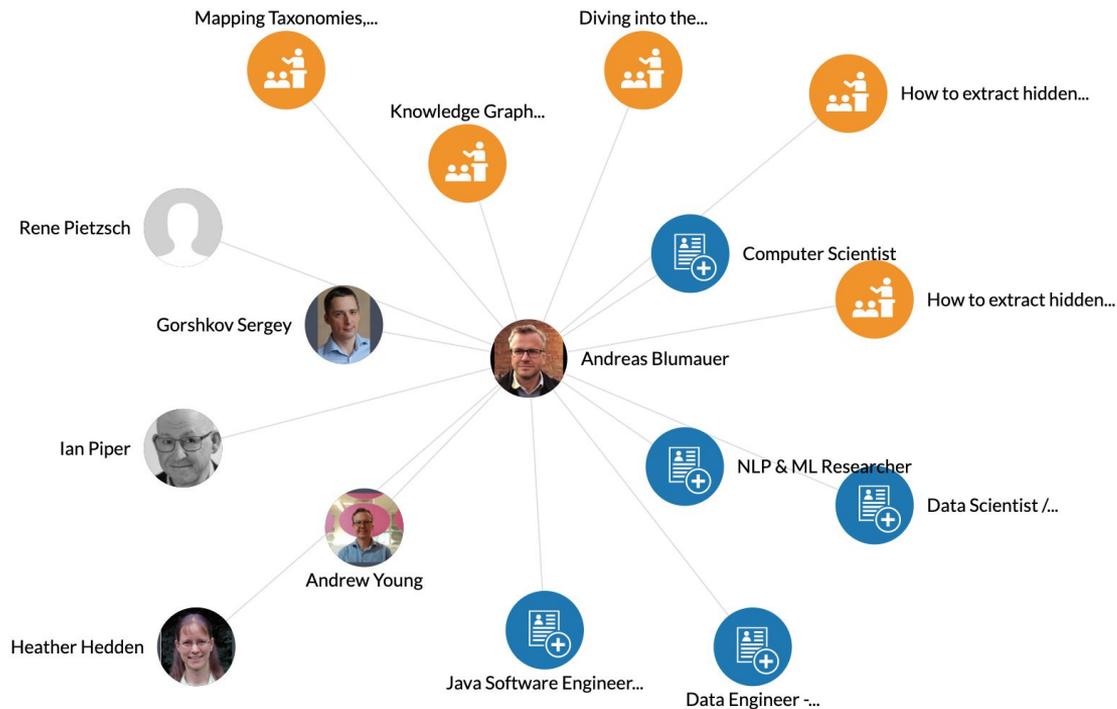
Semantics at Play: Electronic Arts' Linked Data Journey

Knowledge Graphs for **Virtual Assistants**

Integrating Semantics into Dialog Workflows



Use Case Semantic Matchmaking

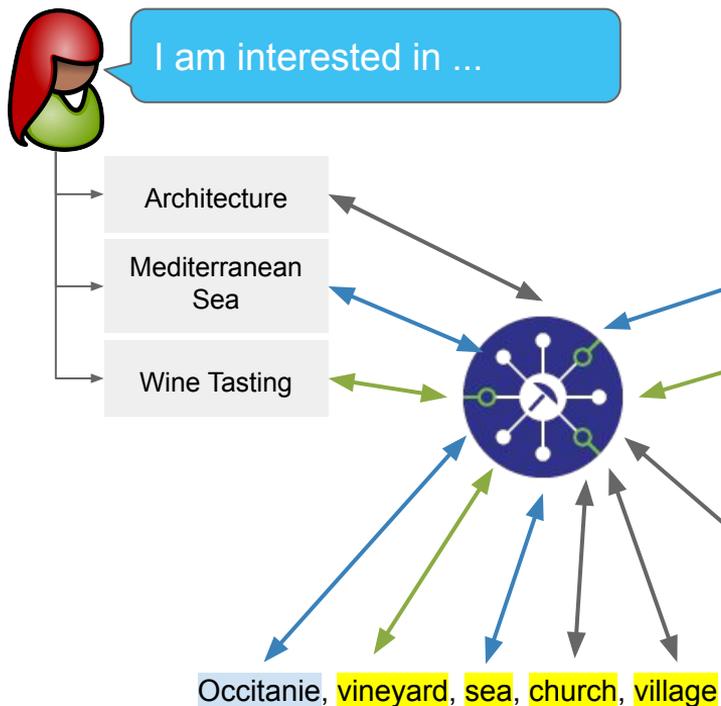


[Event Advisor at SEMANTiCS](#)

Use Case Recommender System

Connecting

- ▶ content to content,
- ▶ people to content,
- ▶ people to people.



Languedoc-Roussillon

Languedoc is a significant producer of wine, and a major contributor to the surplus known as the "wine lake". Today it produces more than a third of the grapes in France, and is a focus for outside investors.

The region contains the historic cities of Carcassonne, Toulouse, Montpellier, countless Roman monuments, medieval abbeys, Romanesque churches, and old castles.

Filter by Topic

- Natural Waters**
 - Coastal Waters
 - Groundwater
 - Lakes
 - Marine Waters
 - Rain and Precipitation
 - Rivers
 - Scarcity, Floods, Droughts (Extreme Events)
 - Transitional Waters
 - Wetlands
- Sustainability**
 - Climate Change
 - Ecosystem Services
 - Limnology and Fresh Water Ecology
 - Nutrient Removal and Recovery
- Water Policy and Administration**
 - Financing
 - IWRM Integrated Water Resource Management

People list view map view

- Sort along your own profile
- [John Smith](#)
Matchmaking score: 622.37
We provide an environmental data hosting solution for a range of companies wanting to manage their water networks. Introducing eagle.io, the...

[view profile](#)

[John Smith](#)
Matchmaking score: 448.18
Sparking, delivering and supporting new ideas and innovations excites me. I am an exceptional communicator, who can bring creativity, leadership and...

[view profile](#)

[John Smith](#)
Matchmaking score: 475.84
I am an expert in International Projects, including the European frameworks as Horizon2020. For CENIEH I set up projects in its diversification areas...

[view profile](#)

[John Smith](#)
Matchmaking score: 386.51
John Peter Smith, PhD is principal scientist at the KWR Watercycle Research Institute and Professor in Water Management and Urban...

[view profile](#)

Example EIP on Water

At the center of the marketplace is the matchmaking function

<https://www.eip-water.eu/>

Keynote

How semantic technologies can support us in making better business decisions

Keynote III

 Time: Wednesday, September 11, 2019 - 09:00 to 10:00



Talks

High-grading Business Decisions through Semantic Technology

Royal Dutch Shell is a knowledge intensive organization. From deep technical expertise in building multi-billion dollar offshore assets in our Upstream Business to driving efficiency and replication in our Retail stations, knowledge is key to our success. The challenge however is finding the right information in the mountain of data we produce on a daily basis. Join us as we take you along the journey of Knowledge Management in Shell over the last decade, and explore together how semantic technologies can support us in making better business decisions.



Andy Boyd



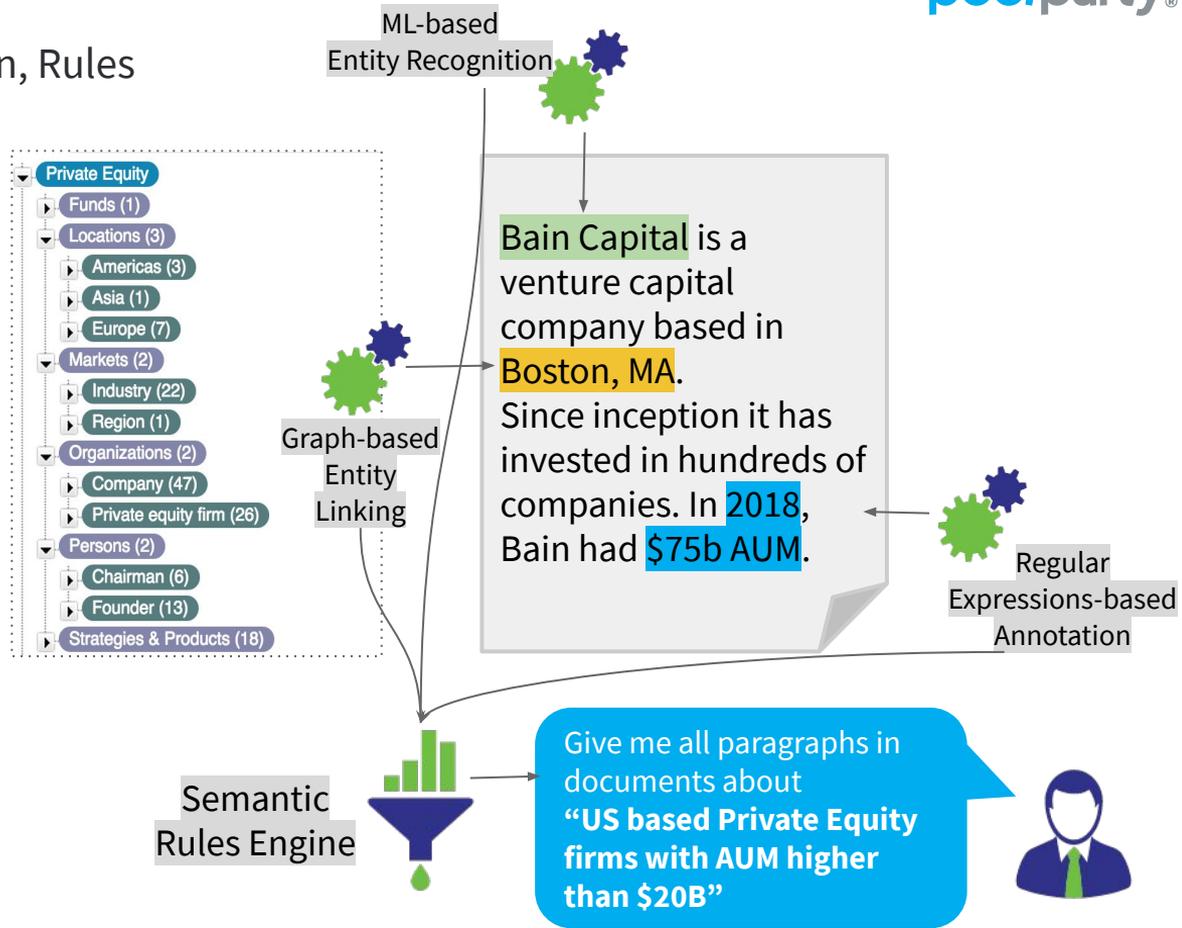
Brendan Nielsen

<https://2019.semantics.cc/keynote-iii-0>

Deep Text Analytics

Annotation, Extraction, Classification, Rules

- ▶ Corpus statistics / Word embeddings
→ **Keyphrase extraction**
- ▶ Graph-based annotation
→ **Entity/Concept linking**
- ▶ Corpus Statistics embedded in graphs
→ **Shadow Concepts**
- ▶ Machine-learning-based annotation
→ **Named entity recognition (NER)**
- ▶ Machine-learning based classification
→ **Document Classification**
- ▶ Annotation based on rules
→ **Regular expressions**



Use Case Contract Intelligence

Semantic analysis and
pre-selection of relevant
clauses in contracts

T CONTRACT AT A GLANCE

Contract facets

▼ 01. Introduction provisions

▼ Parties

[TheCompany](#) 12

[Client](#) 3

▶ 02. Performance and Acceptance

▶ 03. Financial Conditions

[Interest on Late Payment](#) 2

[Pricing revision related to Cola](#) 0

[Pricing revision related to
Currencies rates](#) 0

[T&M based](#) 0

▼ 04. Warranties and Remedies of Performance

[Service warranty](#) 4

[Financial warranty](#) 0

▶ 05. Licences & IPR and Related Warranties and Remedies

▶ 06. Confidentiality - Data Protection - Security

▶ 07. Non-Core obligations

▶ 08. Liability and Insurance

▶ 09. Subcontracting and Transfer of contract

▼ 10. Termination of Contract

[Termination by the Client for
Convenience](#) 8

[Termination by TheCompany for
Convenience](#) 3

[Indefinite term](#) 3

▶ 11. Interpretation of Contract and Disputes

Provided that (1) (2017) received a warning notice to pay thirty (30) days after the date of payment and that the Warning Committee is made aware of (1) (2017) payment delay, late interest shall be charged to (1) (2017) calculated based on the pro-rata interests of a year of 367 days, at a rate equal to three (3) times the French legal rate, and shall be due as from the date when the relevant payment was due pursuant to the Master Service Agreement to the date when such payment shall have been made in full. (1) (2)

(1) (2) is to the extent that as (1) (2017) affiliate fails to carry out, observe or perform any of its payment obligations under the Master Service Agreement, for whatsoever reasons, and that payment remains outstanding 90 (ninety) days after the date on which the Parties had to find an agreement after an amicable procedure, (1) (2017) shall be entitled to terminate the LSA without performing the Reverse Transfer Phase.

In that case, provided that (1) (2017) affiliate received a warning notice to pay thirty (30) days after the date of payment and that the Warning Committee is made aware of (1) (2017) affiliate's payment delay, late interest shall be charged to (1) (2017) affiliate calculated based on the pro-rata interests of a year of 367 days, at a rate equal to three (3) times the French legal rate, and shall be due as from the date when the relevant payment was due pursuant to the Master Service Agreement to the date when such payment shall have been made in full. (1) (2)

Penalties payable by (1) (2017) or as (1) (2017) affiliate constitute a lump sum, fully discharge compensation of (1) (2017) (2017).

(1) (2) may, at its sole discretion, request that any of the obligations of (1) (2017) shall be limited at the end of the month in which they are carried out and shall be paid to (1) (2017) (2017) within the conditions set forth in LSA.

ARTICLE 24. TERM AND TERMINATION

24.1. Term

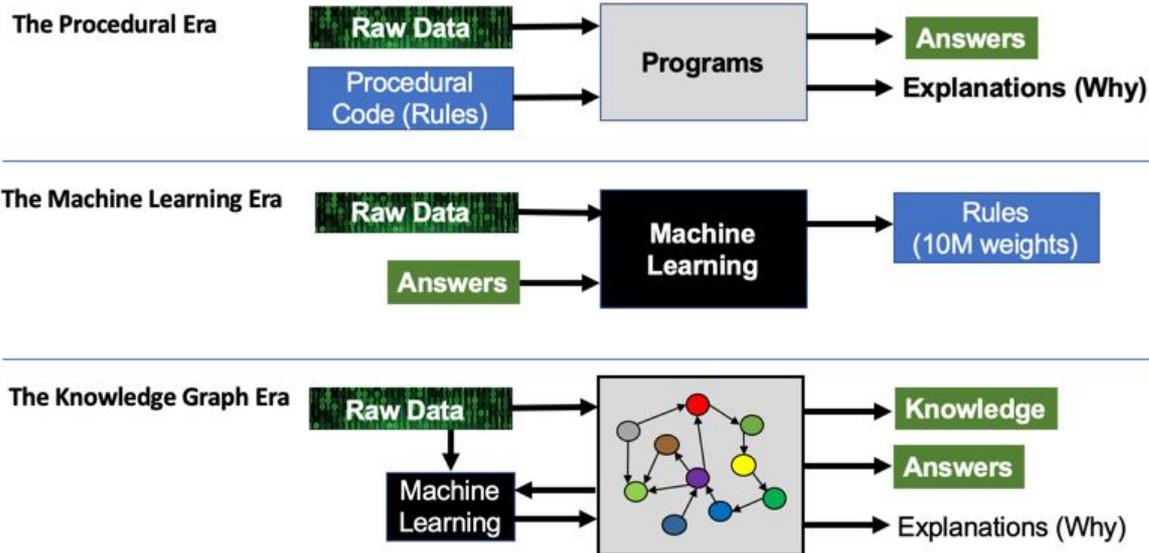
This Master Service Agreement is entered into for a term of five (5) years and five (5) months from the date of signature (hereafter "Initial Term").

Each LSA is signed for a term running from the date of signature until the end of the Initial Term of the MSA, unless the concerned (1) (2017) affiliate shall cease activity.

Unless (1) (2017) affiliate ceases its Master Services, it is renewed at the end of the Initial Term.

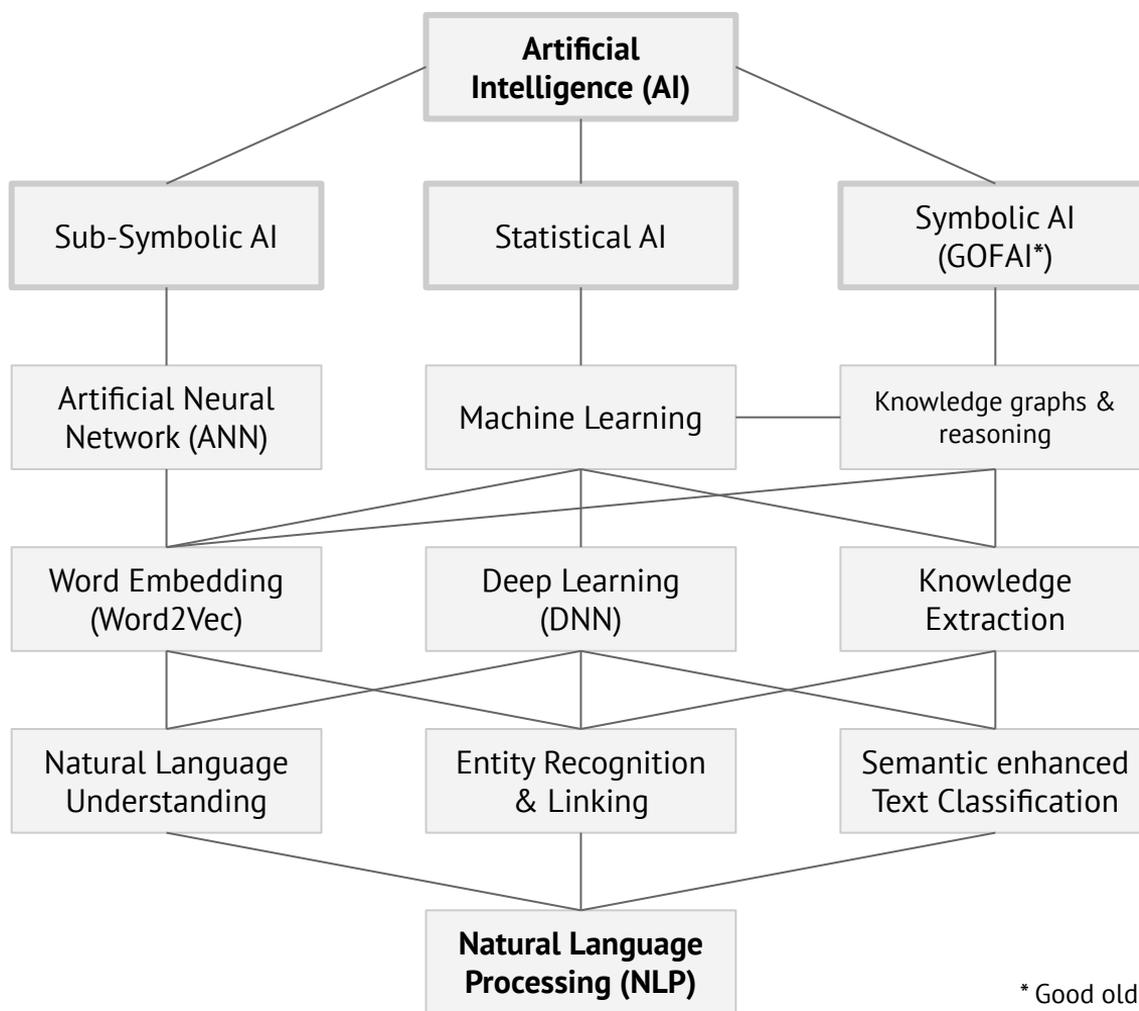
Semantic AI

Towards an Explainable AI.
Fusing Knowledge Graphs,
Machine Learning and NLP



Knowledge Graphs: The Third Era of Computing (Dan McCreary)

Fusing Symbolic & Statistical AI



* Good old-fashioned AI

Semantic Classifier

ML based on semantically enriched training data

The screenshot displays the PoolParty Semantic Classifier interface. On the left, a navigation tree shows 'Training Buckets' (Politics and conflicts, Crime and law, Disasters and accidents) and 'Classifiers' (news channel, news channel-cs2, news channel-cs, news channel-c). The 'news channel' classifier is selected, showing its configuration. The 'Classifier Configuration' tab is active, displaying a dropdown menu for the classifier type (Logistic Regression, Linear Support Vector Machine, Decision Tree, Gradient Boosted Tree, Deep Learning (MLP), Naive Bayes, Random Forest). The 'Features' section includes 'Terms' (checked), 'Concepts', and 'Shadow_Concepts'. The 'Labels' section shows 'Used Labels' (politics, disasters, crime) with checkboxes. The 'Calculated' section displays performance metrics: Last calculation date (10/10/2017 16:06), Performance (f1) (88.78%), Number of training documents (692), and Cross Validation (f1 / recall / precision) (72.55% (f1) / 72.93% (r) / 73.07% (p)). Buttons for 'Train', 'Delete', 'Upload Documents', and 'Duplicate' are visible at the bottom right.

[PoolParty Semantic Classifier](#) combines machine learning algorithms (SVM, Deep Learning, Naive Bayes, etc.) with Semantic Knowledge Graphs.

Use Case Synthetic data generation

Use KGs to generate test data for your ML (which are GDPR compliant)

Abigail Singh
Senior Technical Manager
 6806 disputed rd, M8I 3G8 Belmont
 Canada
 Phone: 882-206-1618
 E-mail: abigail.singh@example.com



About Me

I graduated from Southeast University with a degree in Statistics. I work for United Health as a Senior Technical Manager. I am experienced in the following areas: Expert system, Optical character recognition software, Natural language processing, Machine learning algorithms, Neural network, Speech recognition, and Knowledge base. I am good at decision making, planning and dealing with difficult situations. I see myself as a networker and a problem solver. Colleagues have defined me in the past as a multitasker, logical thinker and strategic planner. I have outstanding proper business etiquette.

My Footprint

Natural language processing, Speech recognition, Similitude, Computer vision, Deep learning, Machine learning algorithms, Optical character recognition software, Collective knowledge, Knowledge base, Expert system, Frontend Developer, Community Manager, Conference Coordinator, Senior Director Product Management

Adam Thomsen
Council Member
 134 østparken, 66926 Roedovre
 Denmark
 Phone: 97948971
 E-mail: adam.thomsen@example.com



About Me

I graduated from Pontificia Universidad Catolica de Chile with a degree in Theoretical Molecular. I currently work as a Council Member at BMW. I am experienced in the following areas: housing operations, Hadoop, Data governance, NoSQL, Indexing, Database systems, and Apply information security policies. I can contribute to the company leadership and presentation skills. The words that best define me are creative, detail oriented, and able to listen. I am good at persuading people, critical thinking and supervising. I am a reliable team builder.

nt

47 Patos de minas

ple.com



Administration at Hunan University. I currently work as a Automation Engineer. I am experienced in the following areas: Social media management, Customer relationship management, Network marketing, Market analysis, and Disseminate corporate information. The words that best defined me in the past as a team player, multitasker and critical thinker.

My Footprint

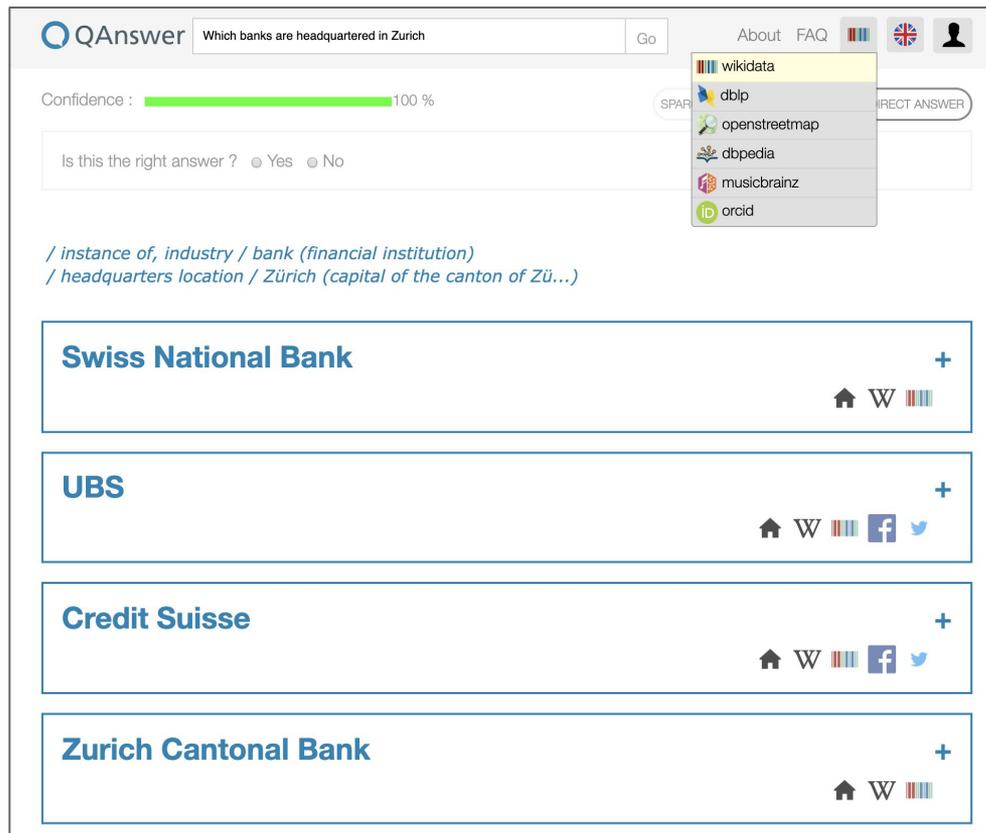
Telemarketing, Pay-per-click, Social content, Marketing department processes, Professional writing, Marketing principles, Social network, Disseminate corporate information, Content marketing strategy, Use online communication tools, Web Developer, PR Manager, Vice President Product Management, Actor

Demo

QA based on KGs

[QAnswer](#) is a so called "knowledge (or ontology) based" QA system

Which banks are headquartered in Zurich?



The screenshot shows the QAnswer web interface. At the top, there is a search bar with the query "Which banks are headquartered in Zurich?" and a "Go" button. To the right of the search bar are links for "About" and "FAQ", a language selector (currently set to English), and a user profile icon. Below the search bar, a "Confidence" bar is shown at 100%. A dropdown menu is open, listing various knowledge graph sources: wikidata (highlighted), dblp, openstreetmap, dbpedia, musicbrainz, and orcid. Below the dropdown, there are two SPARQL query snippets: `/ instance of, industry / bank (financial institution)` and `/ headquarters location / Zürich (capital of the canton of Zü...)`. The main content area displays four search results, each in a blue-bordered box:

- Swiss National Bank**: Includes a home icon, a 'W' icon, and a Swiss flag icon.
- UBS**: Includes a home icon, a 'W' icon, a Swiss flag icon, a Facebook icon, and a Twitter icon.
- Credit Suisse**: Includes a home icon, a 'W' icon, a Swiss flag icon, a Facebook icon, and a Twitter icon.
- Zurich Cantonal Bank**: Includes a home icon, a 'W' icon, and a Swiss flag icon.

Types of solutions based on Knowledge Graphs

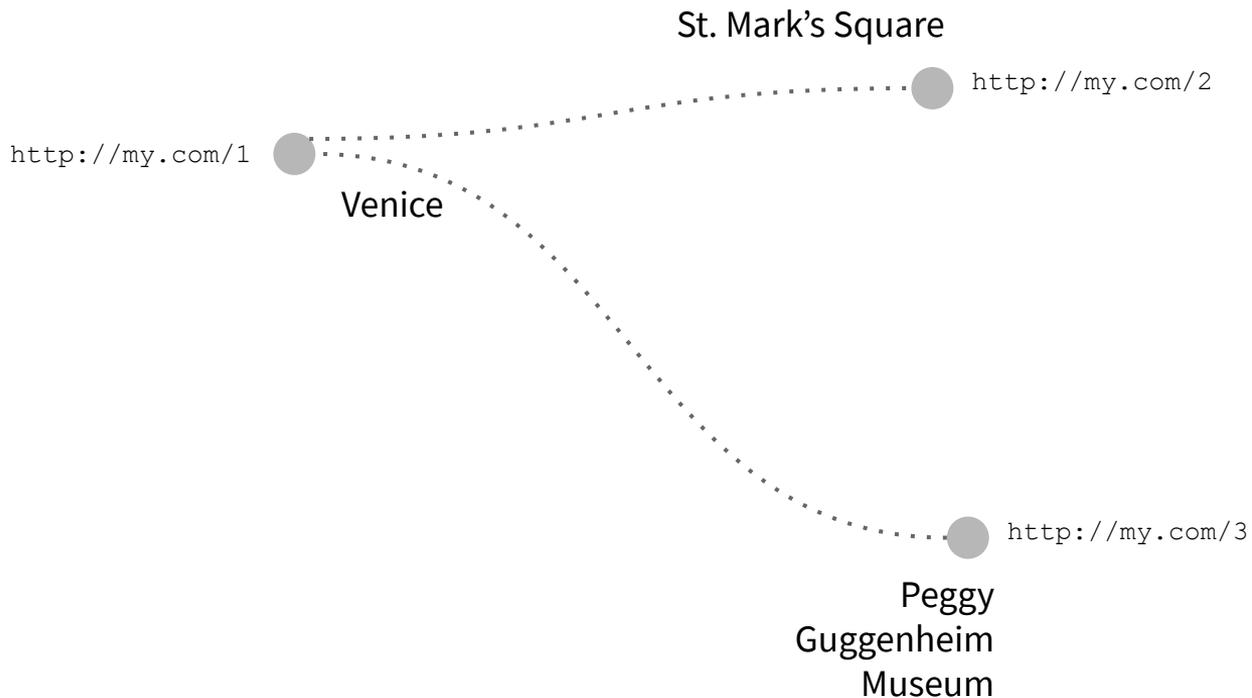
	Data	Content	Knowledge
Front-end	 Process Automation <ul style="list-style-type: none">▶ Enhanced Machine Learning▶ Text Mining & NLP▶ Data Classification▶ Document & Text Validation	 Customer Experience <ul style="list-style-type: none">▶ Recommender Systems▶ SEO▶ Smart Helpdesk Solutions▶ Virtual assistants/ Q&A engines	 Knowledge Management <ul style="list-style-type: none">▶ Graph Analytics▶ Personalization/Skills Mgmt.▶ Knowledge Discovery Portals▶ Semantic Search
Back-end	 Agile Data Integration <ul style="list-style-type: none">▶ Linked Data▶ Integrating heterogeneous data▶ Entity Linking	 Information Management <ul style="list-style-type: none">▶ Semantic Content Management▶ Metadata Management▶ Masterdata Management	 Knowledge Engineering <ul style="list-style-type: none">▶ Taxonomy Management▶ Ontology Management▶ Knowledge Graph Mgmt.▶ Data Visualization

THINGS, NOT STRINGS

From Taxonomies over Ontologies to Knowledge Graphs

Things and URIs

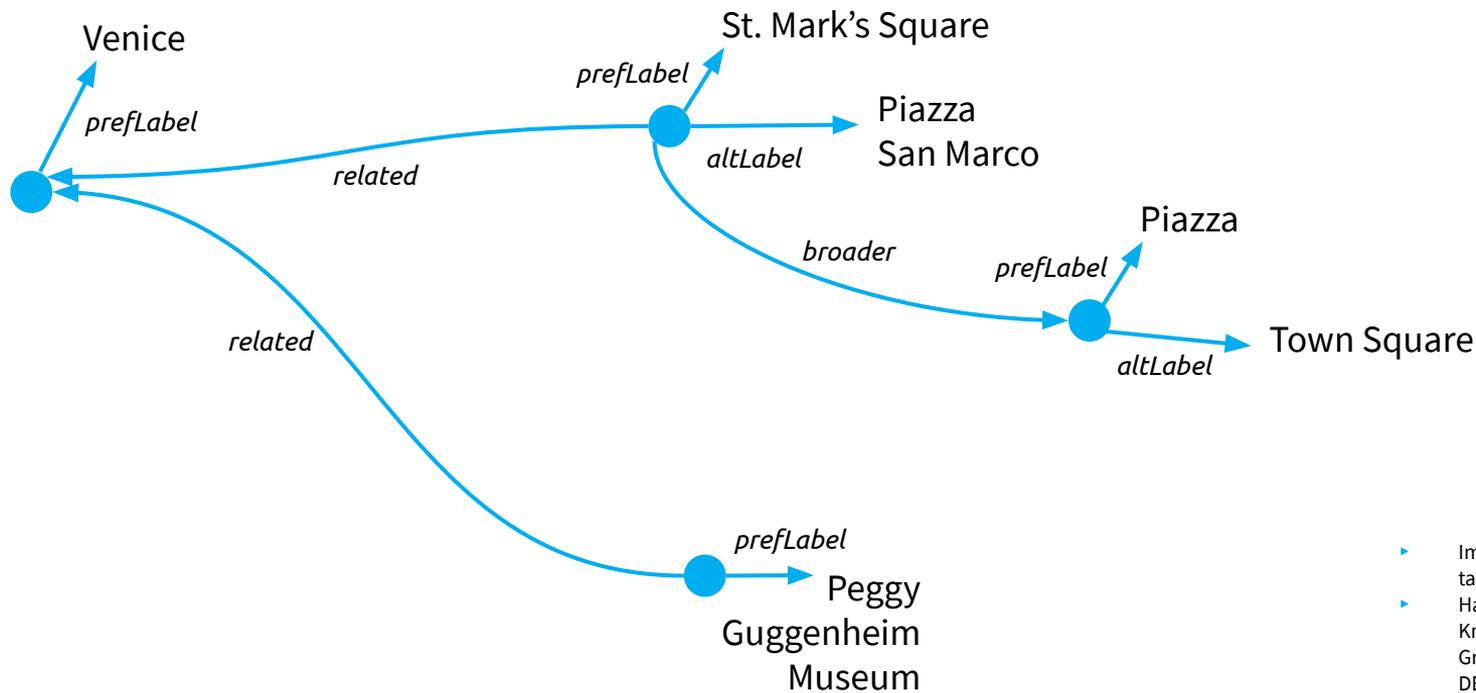
This is not yet a Knowledge Graph!



- ▶ Learn from text corpora
- ▶ Co-occurrences and word embeddings
- ▶ Extract entities via ML
- ▶ Import CSV/Excel

Labels and basic relations

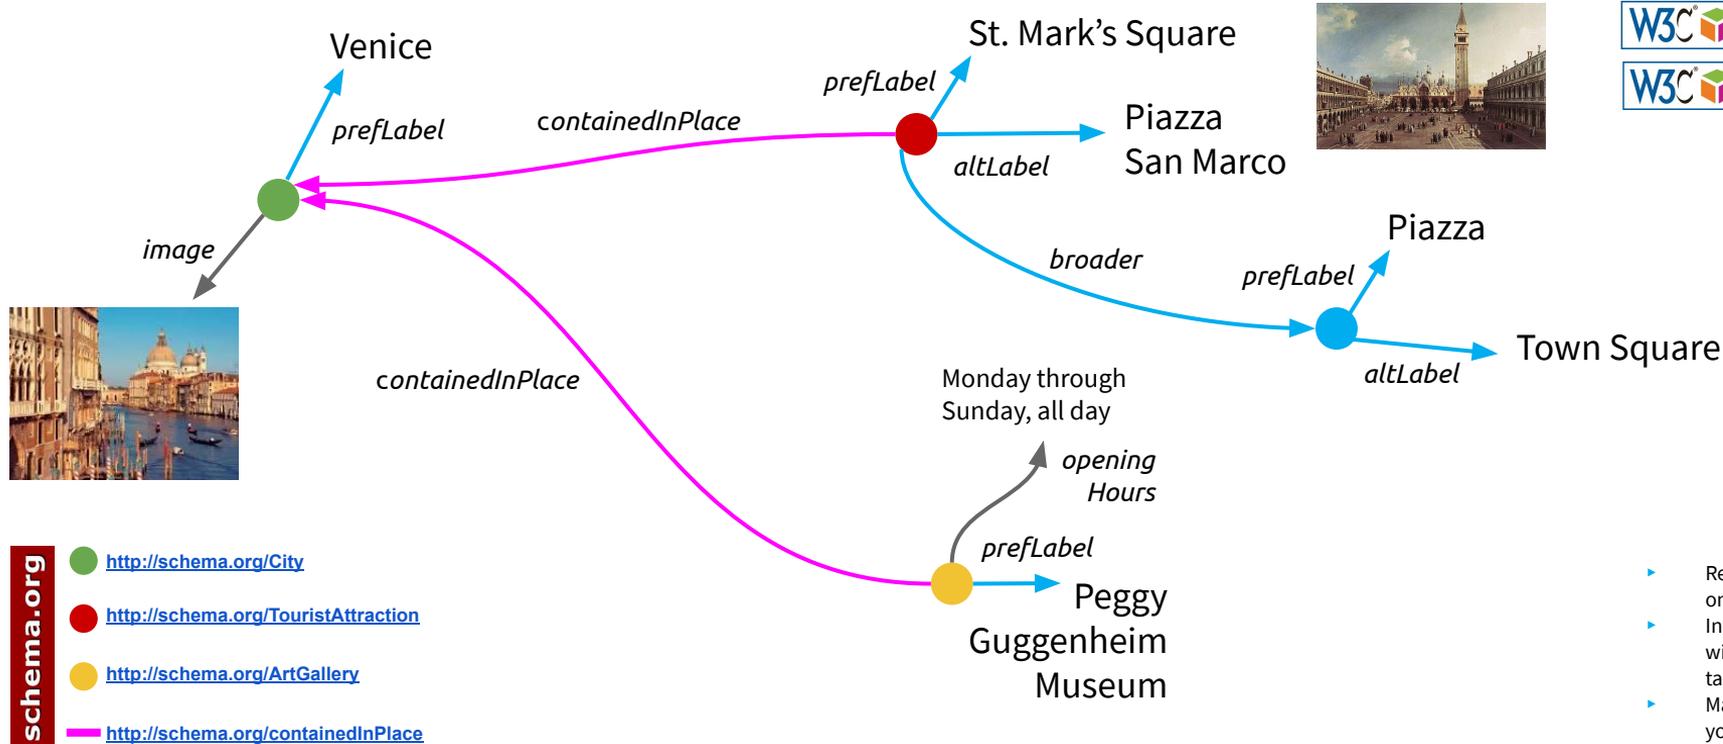
This is the backbone of a Knowledge Graph



- ▶ Import known taxonomies
- ▶ Harvest from Knowledge Graphs like DBpedia
- ▶ Ingest from own DBs/XML

Classes and specific relations

Ontologies give your knowledge graph an additional "dimensionality"



- ▶ Reuse existing ontologies
- ▶ Interlink them with your taxonomies
- ▶ Map them to your DBs/XML

Metadata and Graph annotations

Enrich and qualify data in your knowledge graph



CC BY-SA 3.0

image



Venice

prefLabel

St. Mark's Square

prefLabel

Piazza San Marco

altLabel



Piazza

prefLabel

Town Square

altLabel

Monday through Sunday, all day

opening Hours



Peggy Guggenheim Museum

containedInPlace

containedInPlace

broader

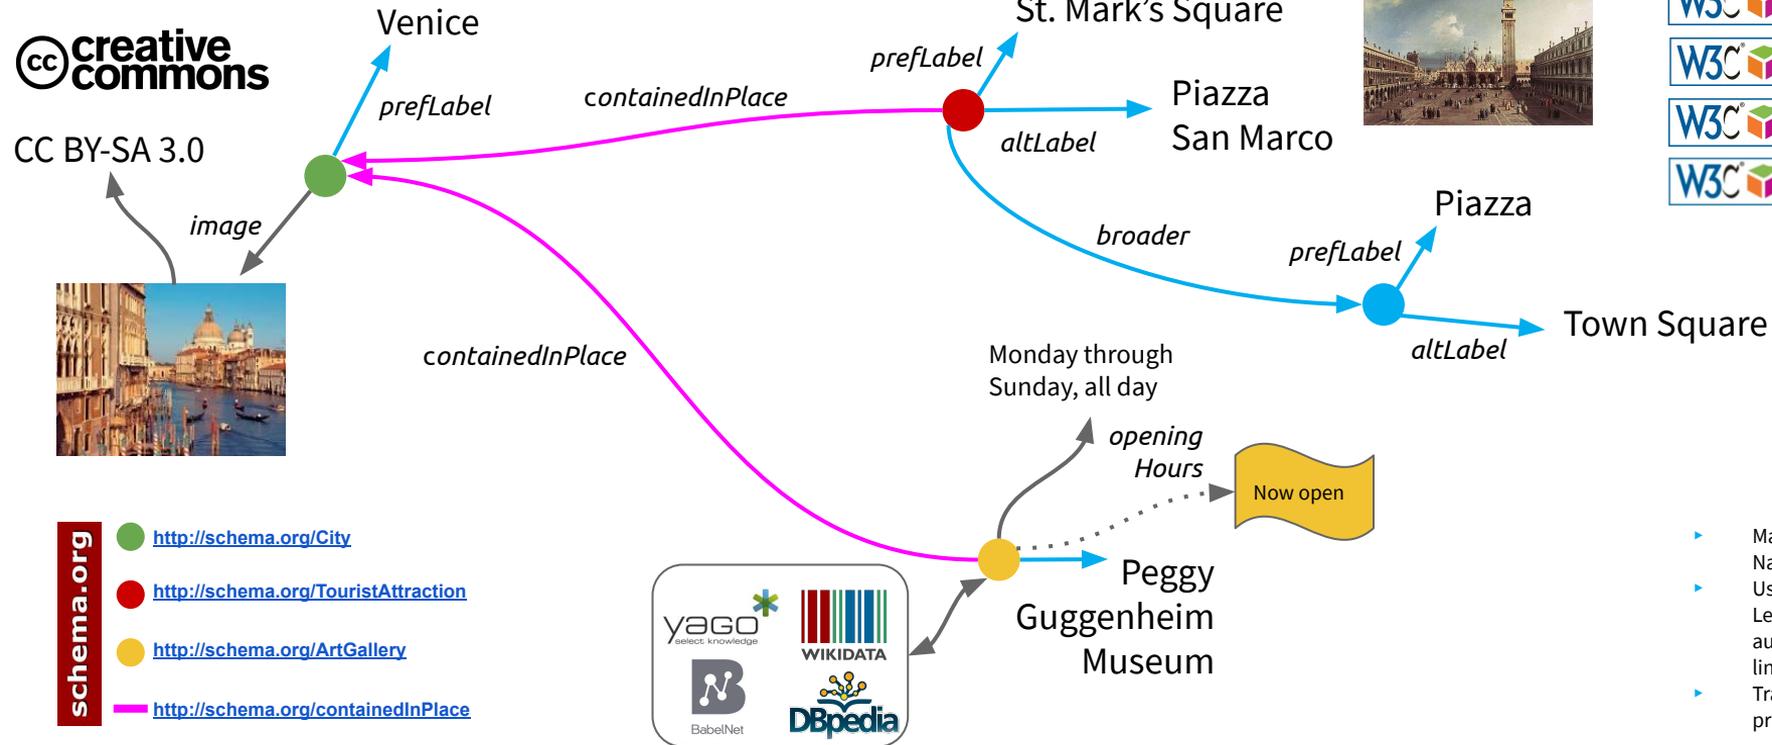
schema.org

- <http://schema.org/City>
- <http://schema.org/TouristAttraction>
- <http://schema.org/ArtGallery>
- <http://schema.org/containedInPlace>

- ▶ Reuse existing schemas
- ▶ Make use of SPARQL and reasoning
- ▶ Make use of constraint languages like SHACL

Graph linking

Additional facts for your knowledge graph, nearly for free!

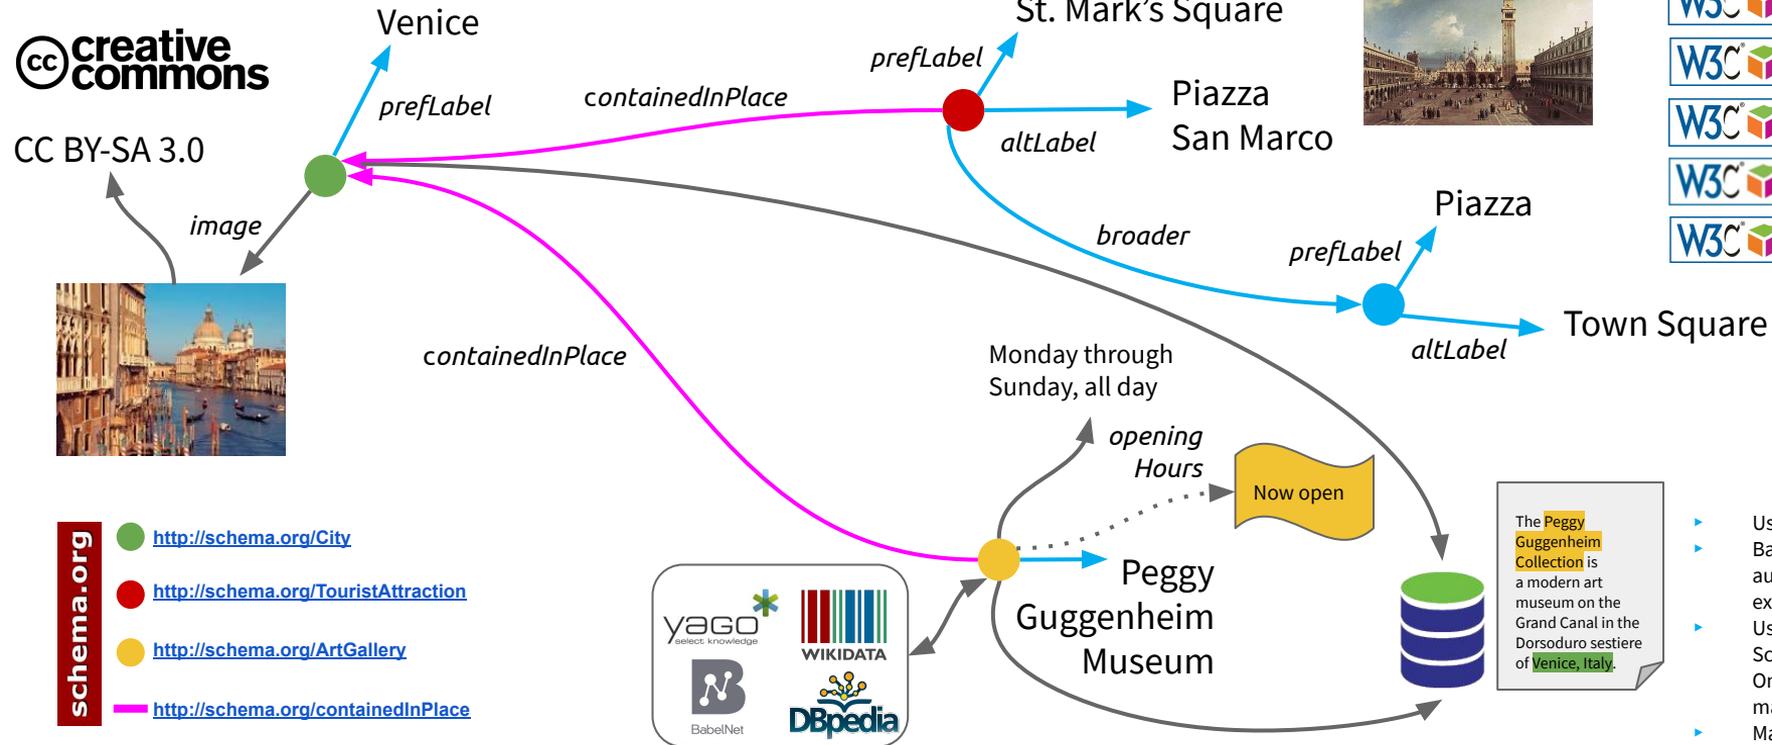


- ▶ Make use of Named Graphs
- ▶ Use Machine Learning for automatic linking
- ▶ Track data provenance

Link all your data!

Bring structured/unstructured enterprise data into the knowledge graph

- SKOS
- OWL2
- RDF
- SPARQL
- SHACL
- R2RML



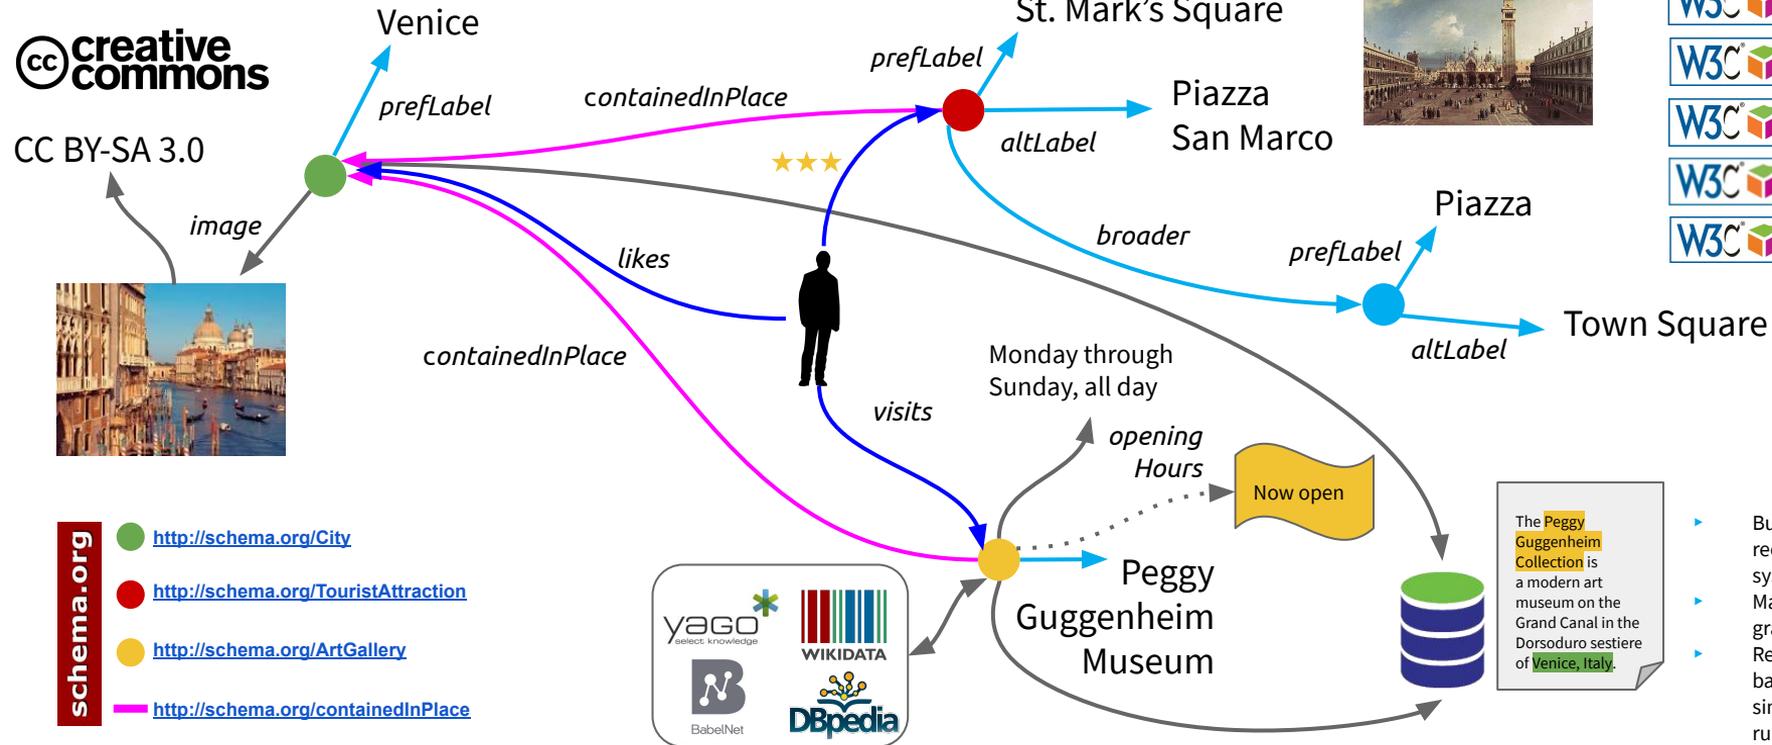
- ▶ Use Text Mining
- ▶ Based on automatic entity extraction
- ▶ Use R2RML Schema to Ontology mapping
- ▶ Make use of ML

- schema.org**
- <http://schema.org/City>
 - <http://schema.org/TouristAttraction>
 - <http://schema.org/ArtGallery>
 - <http://schema.org/containedInPlace>

Putting the user into the graph

Provide all your data personalized and contextualized

- W3C SKOS
- W3C OWL2
- W3C RDF
- W3C SPARQL
- W3C SHACL
- W3C R2RML



- ▶ Build recommender systems
- ▶ Make use of graph analytics
- ▶ Recommender based on similarity and/or rules

Who is **CEO** of a **Bank**,
headquartered in Europe that
generates revenue per employee
higher than 400,000 Euro?



Answers instead of search results



	Head-quarter	Ticker Symbol	Revenue
Credit Suisse	Zurich	VTX: CSGN	CHF 23.4b
HSBC	London	LON: HSBA	USD 60.0b
Allianz	Munich	ETR: ALV	EUR 122.3b
Deutsche Bank	Frankfurt	ETR: DBK	EUR 33.5b



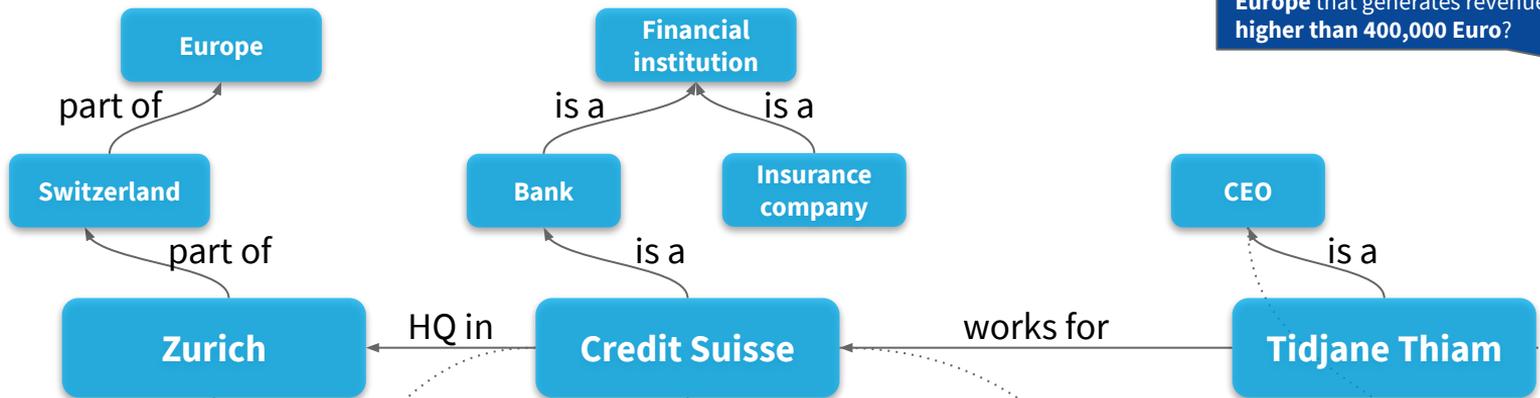
```
<Employees>
<VTX:CSGN>
  48,200
</VTX:CSGN>
<LON:HSBA>
  266,273
</LON:HSBA>
<ETR:ALV>
  147,425
</ETR:ALV>
<ETR:DBK>
  101,104
</ETR:DBK>
</Employees>
```



The task awaiting Tidjane Thiam when he takes over from Brady Dougan as the new chief executive at Credit Suisse Group AG is clear: how to pull the Swiss bank out of a post-financial crisis rut.

KGs: Linking, Mapping, Reasoning

Who is **CEO** of a **Bank**, headquartered in **Europe** that generates revenue per employee higher than 400,000 Euro?



Organisation	HQ	Umsatz
CS	Zürich	CHF 23.4b
HSBC	London	USD 60.0b
Allianz	München	EUR 122.3b
Deutsche Bank	Frankfurt	EUR 33.5b



```

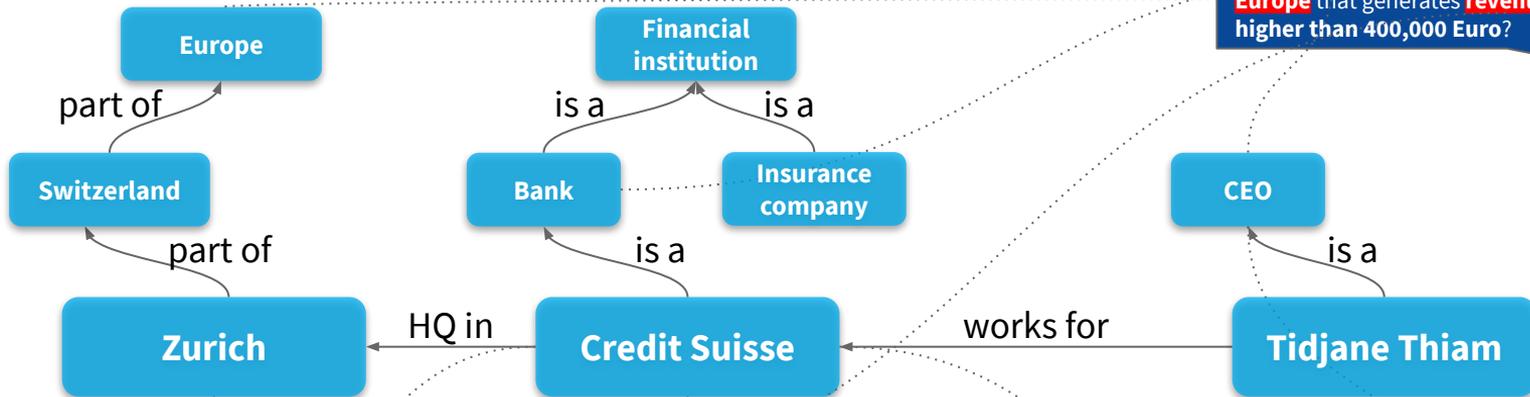
<Employees>
  <VTX:CSGN>48,200</VTX:CSGN>
  <LON:HSBA>266,273</LON:HSBA>
  <ETR: ALV>147,425</ETR: ALV>
  <ETR: DBK>101,104</ETR: DBK>
</Employees>
  
```



The task awaiting **Tidjane Thiam** when he takes over from Brady Dougan as the new **chief executive** at **Credit Suisse Group AG** is clear: how to pull the Swiss bank out of a post-financial crisis rut.

Knowledge Graphs support NLU and QA

Who is **CEO** of a **Bank**, headquartered in **Europe** that generates **revenue** per employee higher than 400,000 Euro?



Organisation	HQ	Umsatz
CS	Zürich	CHF 23.4b
HSBC	London	USD 60.0b
Allianz	München	EUR 122.3b
Deutsche Bank	Frankfurt	EUR 33.5b



```

<Employees>
  <VTX:CSGN>48,200</VTX:CSGN>
  <LON:HSBA>266,273</LON:HSBA>
  <ETR: ALV>147,425</ETR: ALV>
  <ETR: DBK>101,104</ETR: DBK>
</Employees>
  
```

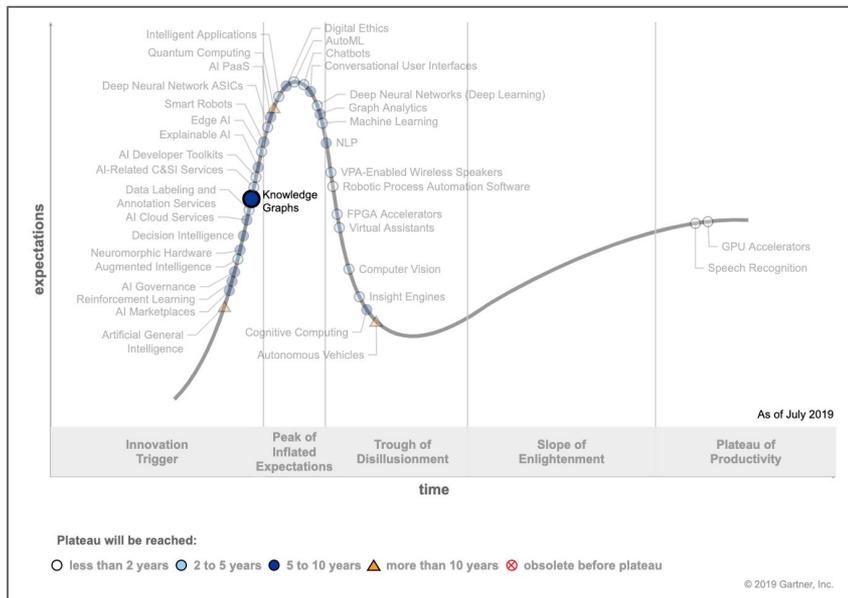


The task awaiting **Tidjane Thiam** when he takes over from Brady Dougan as the new **chief executive** at **Credit Suisse Group AG** is clear: how to pull the Swiss bank out of a post-financial crisis rut.

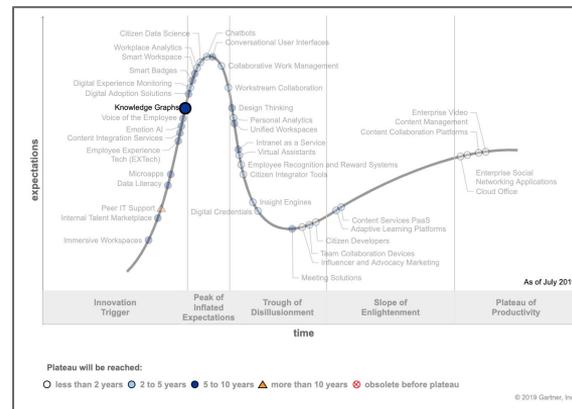
CHANGE MANAGEMENT

How to implement disruptive technologies in organizations?

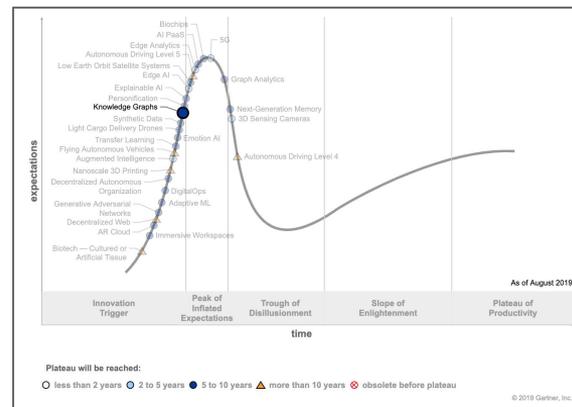
C-Level: Knowledge Graphs in various Hype Cycles



Hype Cycle for Artificial Intelligence



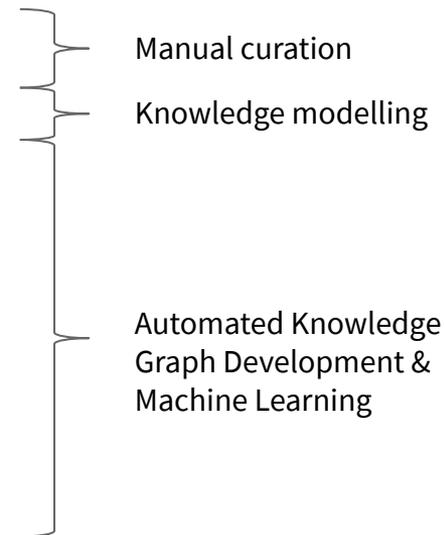
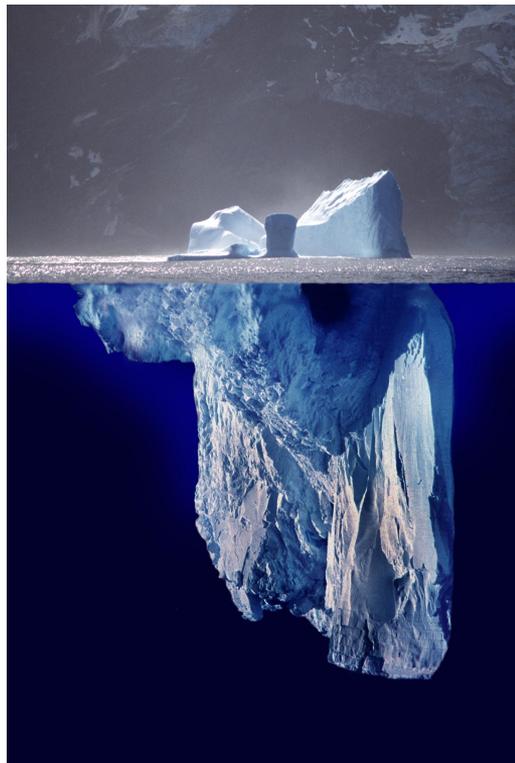
Hype Cycle for the Digital Workplace



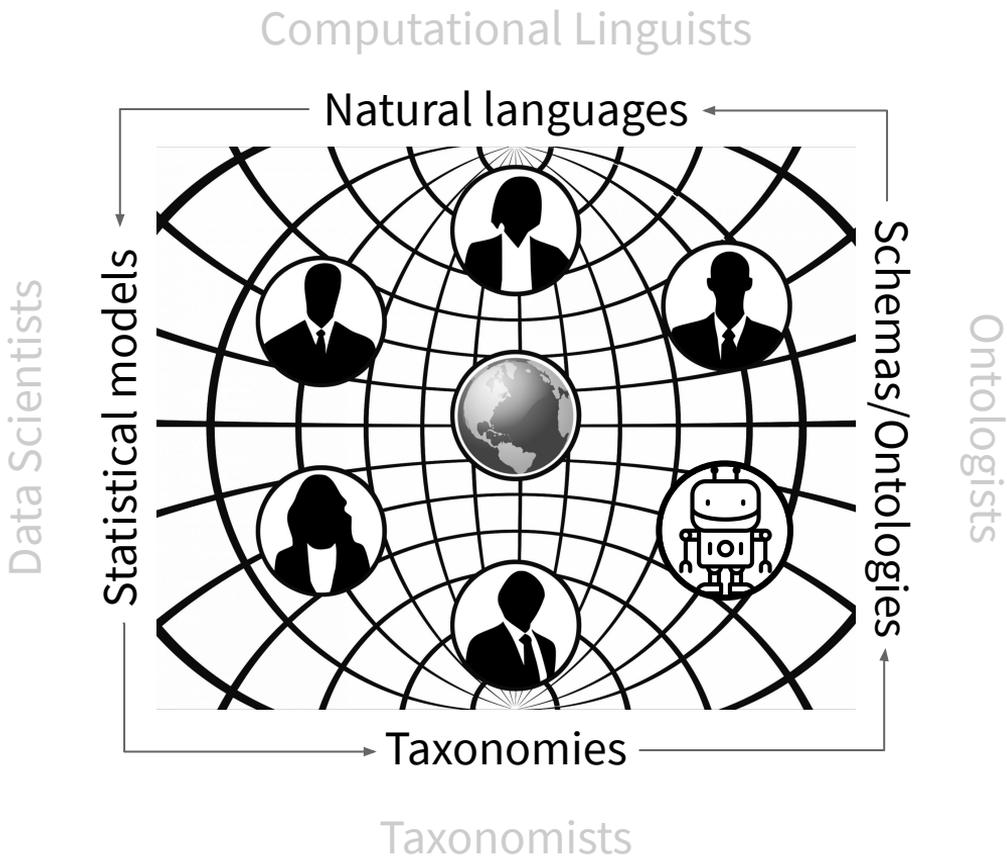
Hype Cycle for Emerging Technologies

Linked Data Life Cycle

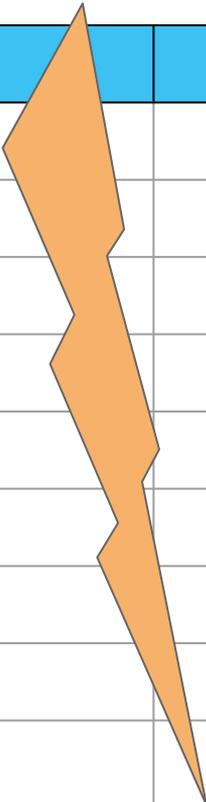
Data Engineers + Subject Matter Experts + Machine Learning working together



Our approach
Semantic AI

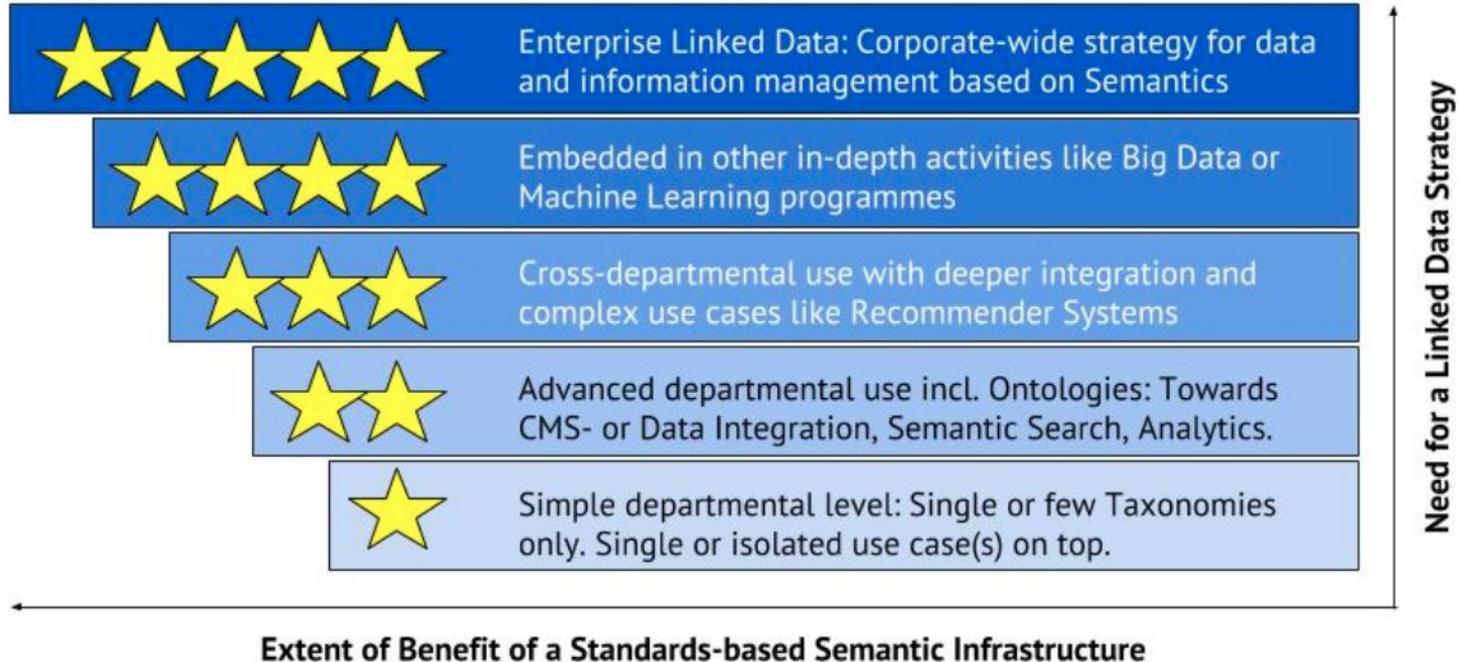


Bringing two types of mindsets together



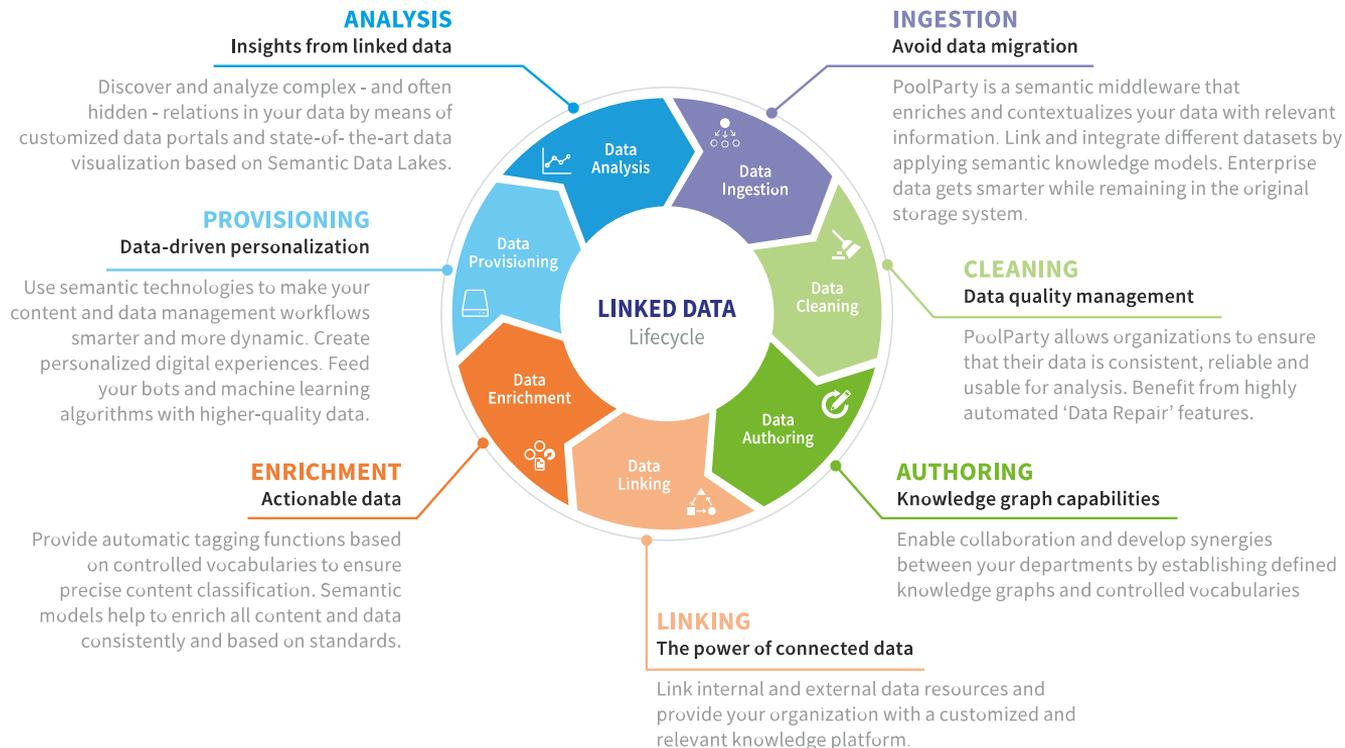
Data-centric people	Knowledge-centric people
Service-oriented industries	Asset-oriented industries
Databases & Excel	Content & documents
Structured data	Unstructured data
Linked Data	Knowledge Graph
Algorithms	Collaborative knowledge management
Predictability and automation	Product innovation and risk mitigation
Ontologies and Machine Learning	Taxonomies and Collaboration
“Actionable Data!”	“Better Decisions!”
Understand the customer and market place better	Stay or become the expert in the field

Enterprise Semantics Maturity Model



Knowledge Graph Management

Data Governance along the Linked Data Life Cycle



Semantic Web Starter Kit

PoolParty Academy

Access to online learning



Prototype Specifications

Technical and functional description of features



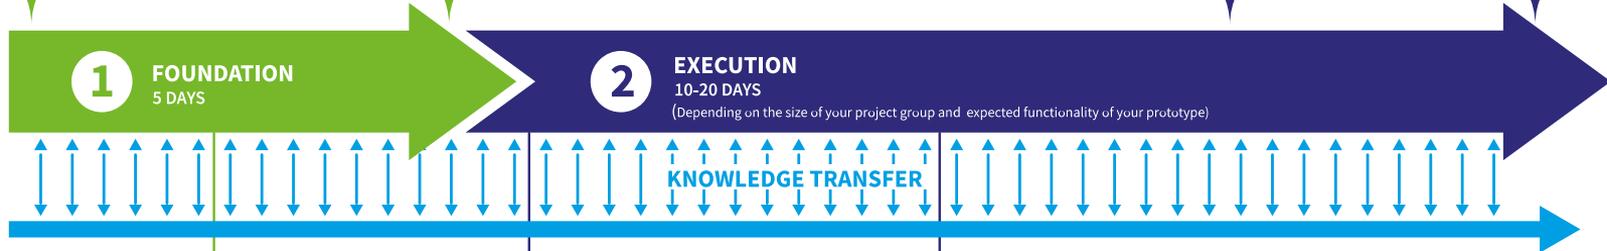
Your Prototype

Functional prototype as defined in requirements



Strategy White Paper

Recommendations and roadmap



Kick-off Workshop

Aligning business & IT strategy



Test environment

Installation of system at customer site or cloud server



Data Analysis & Processing

Review of available data & data processing



[Learn more about it!](#)

Make things work

SEMANTIC AI

Build your own applications

Fact sheet: PoolParty Semantic Suite



Most complete and secure
Semantic Middleware on
the Global Market

 **Semantic AI:**
Fusing Graphs, NLP,
and Machine Learning

ISO 27001:2013 certified
W3C standards compliant

First release in **2009**

Current version **7.0**

On-premise or
cloud-based



Over **200** installations
world-wide



Named as Sample
Vendor in
**Gartner's Hype
Cycle for AI 2018**



Named as
Representative
Vendor in **Gartner's
Market Guide for
Hosted AI Services
2018**



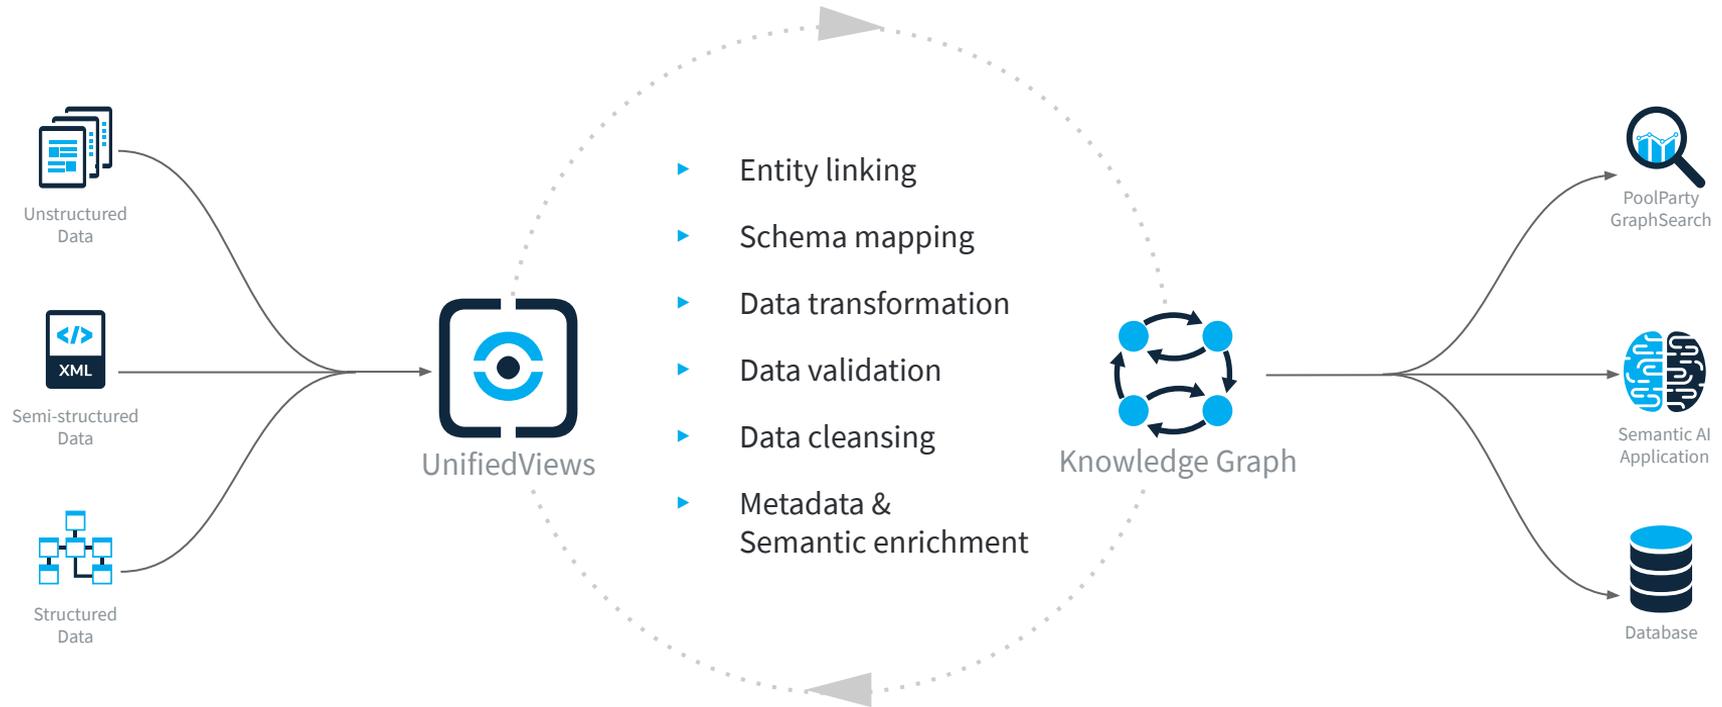
KMWorld listed
PoolParty as
**Trend-Setting
Product** 2015, 2016,
2017, and 2018

“PoolParty brings the power of Google into the enterprise environment through Knowledge Graphs.”

Knowledge Architect at a Top 3 Oil & Gas company

How does it work?

The Data Engineer's perspective



Example: HR Analytics

Demo

As an HR manager, for upcoming training programmes, I want to identify employees who

- have a certain skill set
- have a specific degree
- have skills that are increasingly important on the labour market
- fall into a specific salary range

→ Linking Structured to Unstructured Data

Employee database



Resumes



Labour market statistics

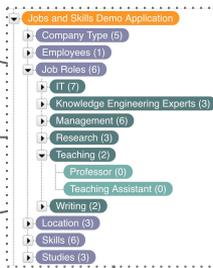


Employee database 

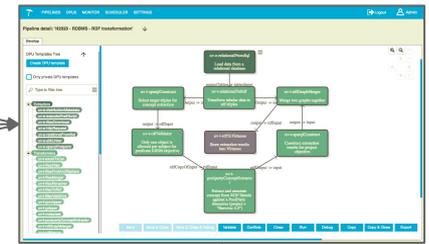
Resumes  SharePoint

Labour market statistics 

PoolParty Thesaurus Server



PoolParty UnifiedViews



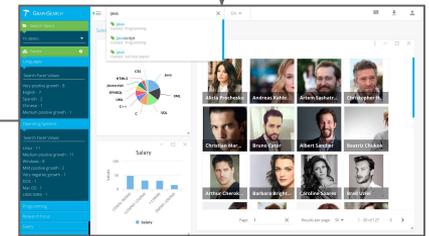
RDF Graph Database

How it works

Now I can identify employees along many dimensions.



PoolParty User

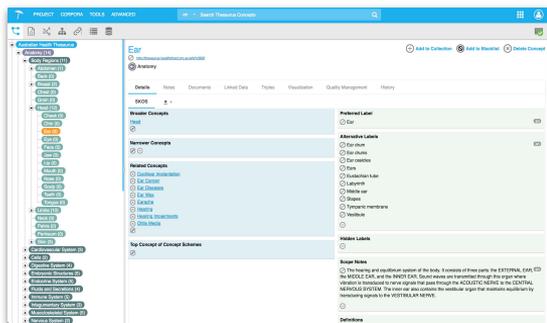


PoolParty GraphSearch

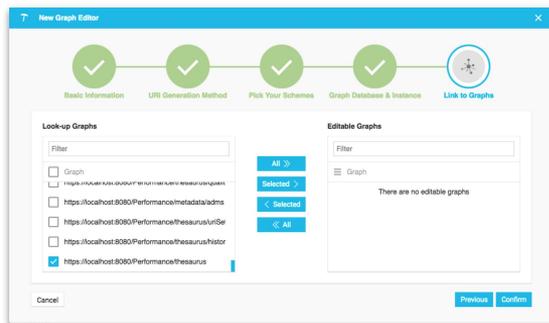
Components and Features



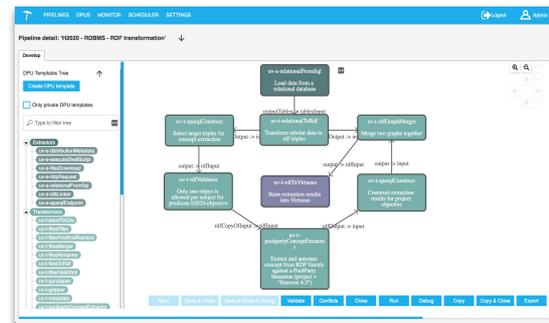
PoolParty Components



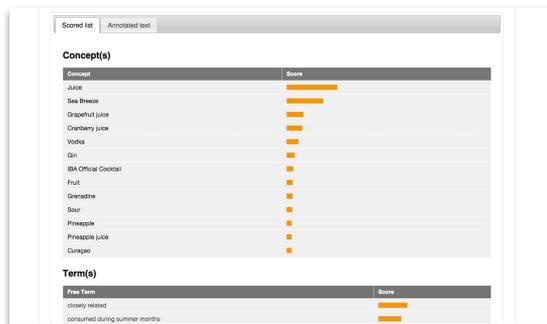
PoolParty Taxonomy & Ontology Server



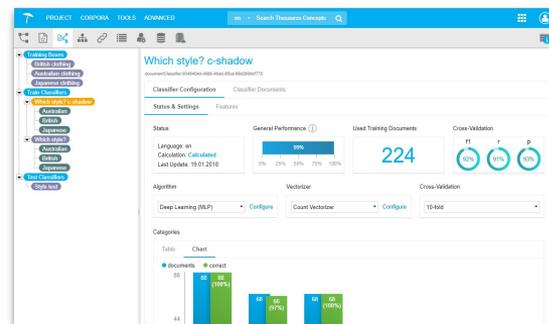
PoolParty GraphEditor



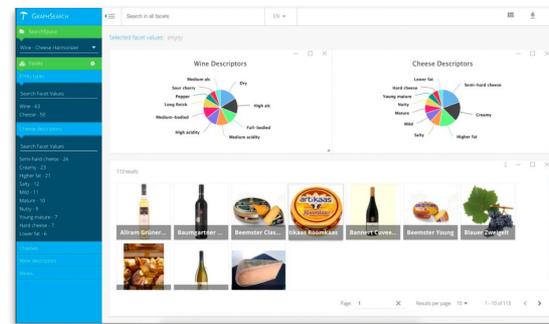
PoolParty UnifiedViews



PoolParty Extractor



PoolParty Semantic Classifier



PoolParty GraphSearch

TECHNICAL DEEP DIVE

Make your data actionable!

Why use Graph Databases?



Dealing with **hierarchical** and **highly connected** datasets



Entity-centric views (in contrast to document-centric views)



Exploring the **connection between entities** of a graph



Integrating heterogeneous data sources
(structured & unstructured in a “schema-late” approach)

What makes a store RDF ready?



It implements at least one of the RDF ready libraries/frameworks:

- ▶ Eclipse **RDF4J**
- ▶ Apache **Jena**

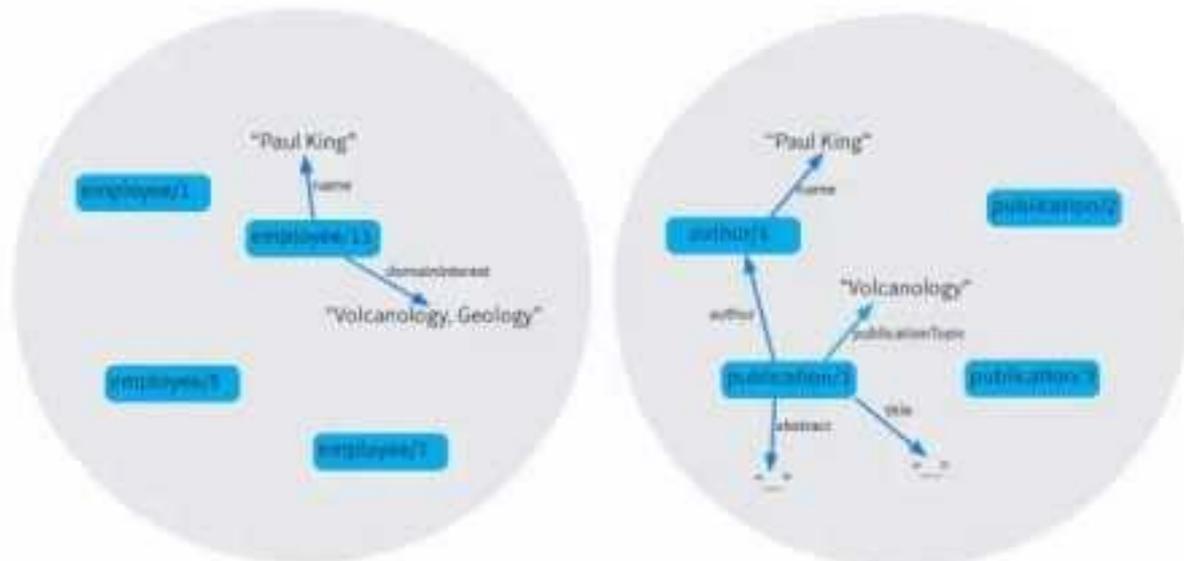


It offers a **SPARQL** endpoint to query the database using SPARQL

SPARQL

SPARQL retrieves and manipulates data stored in an RDF graph database

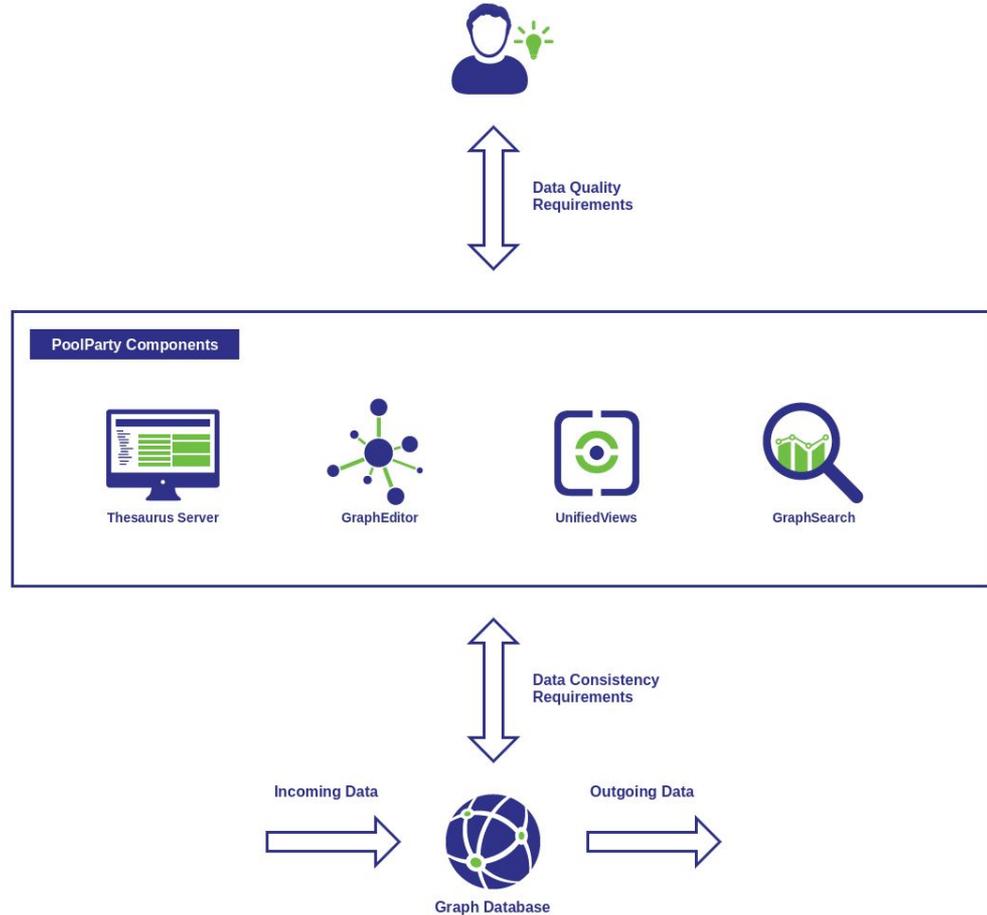
[Watch Tutorial from PoolParty Academy](#)



SHACL

Shapes
Constraint
Language

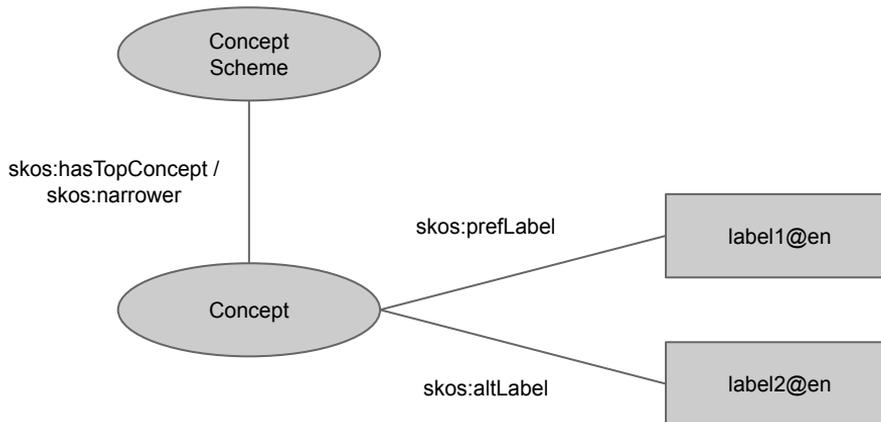
Data Quality and Consistency:
Validating graph-based data
against a set of conditions



Concept labels and hierarchic structure

Example 1

Validating Concept labels and hierarchic structure in a SKOS taxonomy



Constraints:

A concept has to connect to a concept scheme.

The labels have to be disjoint and the language must be unique.

Maximum active board members

Example 2

Compliance checks



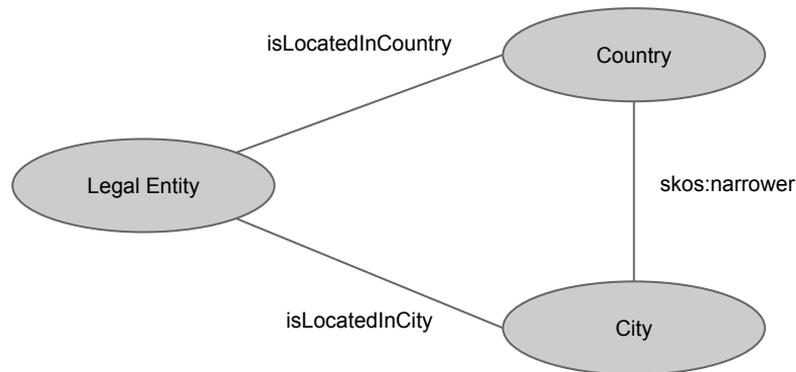
Constraint:

There must not be more than two active board members for each Legal Entity.

Consistent geographic information

Example 3

Data consistency



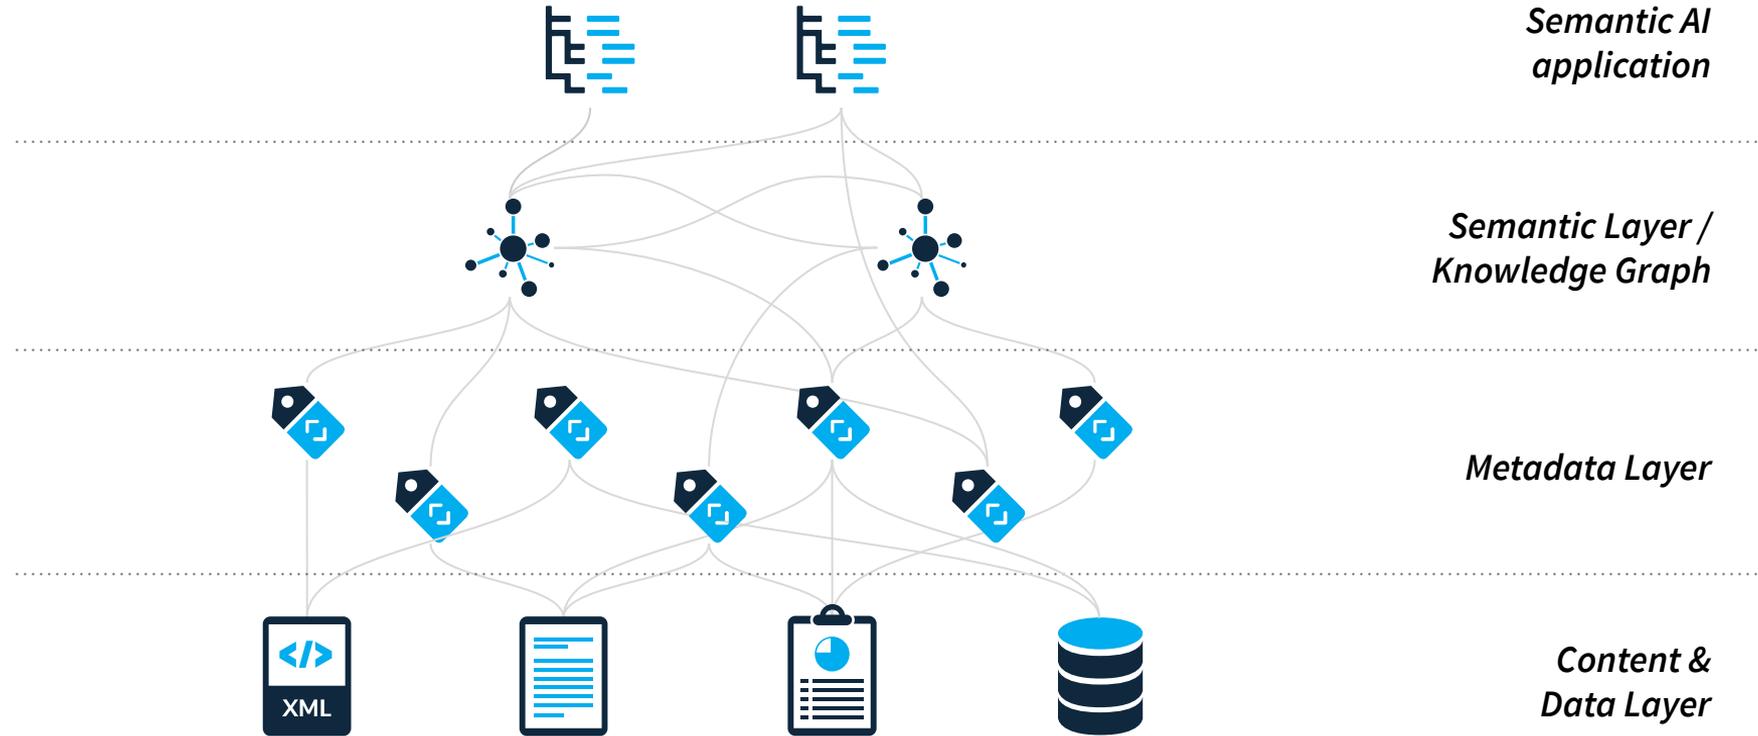
Constraint:

If a Legal Entity has a country and a city assigned, then both must be related with a skos:narrower path, so that the geo information is consistent.

TAKE AWAYS

Make your data actionable!

Four-layered Information Architecture



From Documents To Entities

Watch out for ambiguity!

All
 Images
 News
 Maps
 Videos
 More

Settings Tools

About 739.000.000 results (0,81 seconds)

[Images for jaguar car](#)

[→ More images for jaguar car](#)

[Report images](#)

Market Selector | Jaguar | View the site in your preferred language
<https://www.jaguar.com/market-selector.html> ▼

Discover the different language sites we have to make browsing our vehicle range's easier. We have over 100 different language options available. Learn more.

Jaguar XF · Jaguar XJ | Luxury Saloon Car · Jaguar I-PACE · Jaguar XE

Jaguar Cars - Wikipedia
https://en.wikipedia.org/wiki/Jaguar_Cars ▼

Jaguar is the luxury vehicle brand of Jaguar Land Rover, a British multinational car manufacturer with its headquarters in Whitley, Coventry, England. Jaguar ...

Previous owners: Jaguar Cars (1935–2012) Related brands: Land Rover
 Country: United Kingdom Produced by: Jaguar Land Rover

Jaguar XJ · Jaguar E-Pace · Jaguar XK · Jaguar F-Type

Videos

2016 Jaguar XF First Look - 2015 New York Auto Show

AutoGuide.com
YouTube - Apr 1, 2015

2018 Jaguar XJ - Review

AutoShow
YouTube - Jul 29, 2017

Jaguar XF | Features and Benefits

Jaguar
YouTube - Oct 14, 2016

Jaguar Cars – Wikipedia
https://de.wikipedia.org/wiki/Jaguar_Cars ▼ [Translate this page](#)

Jaguar (engl. Aussprache: [ˈdʒæɡjuə] in Großbritannien, [ˈdʒæɡ.wu] in den Vereinigten Staaten) ist eine Automobil-Marke der Tochterfirma Jaguar Land ...

Jaguar Cars

Luxury vehicles company

Jaguar is the luxury vehicle brand of Jaguar Land Rover, a British multinational car manufacturer with its headquarters in Whitley, Coventry, England. [Wikipedia](#)

Headquarters: Coventry, United Kingdom

Founded: September 4, 1922, Blackpool, United Kingdom

CEO: Ralf Speth (Feb 18, 2010–)

Tagline: "The Art of Performance"

Parent organizations: Tata Motors, British Leyland, British Motor Holdings, Jaguar Land Rover Holdings Limited

Founders: William Lyons, William Walmsley

Profiles

People also search for [View 10+ more](#)

Tata Motors

Aston Martin

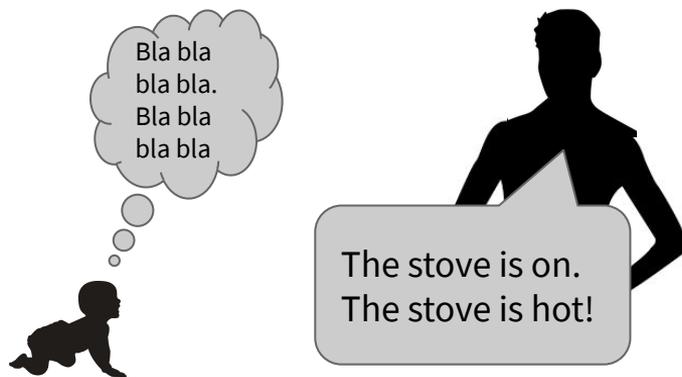
Mercedes-Benz

Bentley Motors Limited

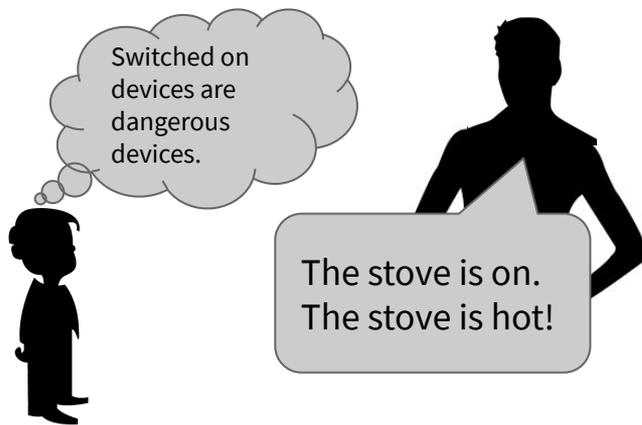
Land Rover

ML based on Semantic AI Engines

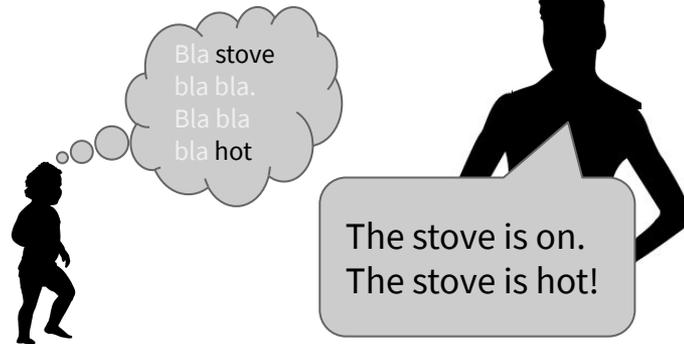
Separating signal from noise



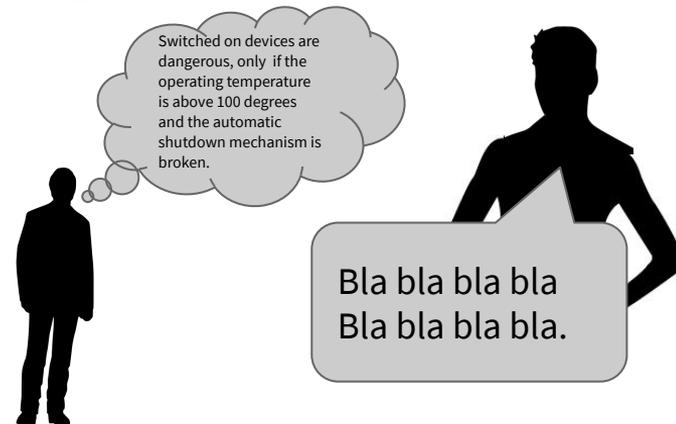
Taxonomical model → is-a abstractions



Word embeddings/cooccurrences → is related

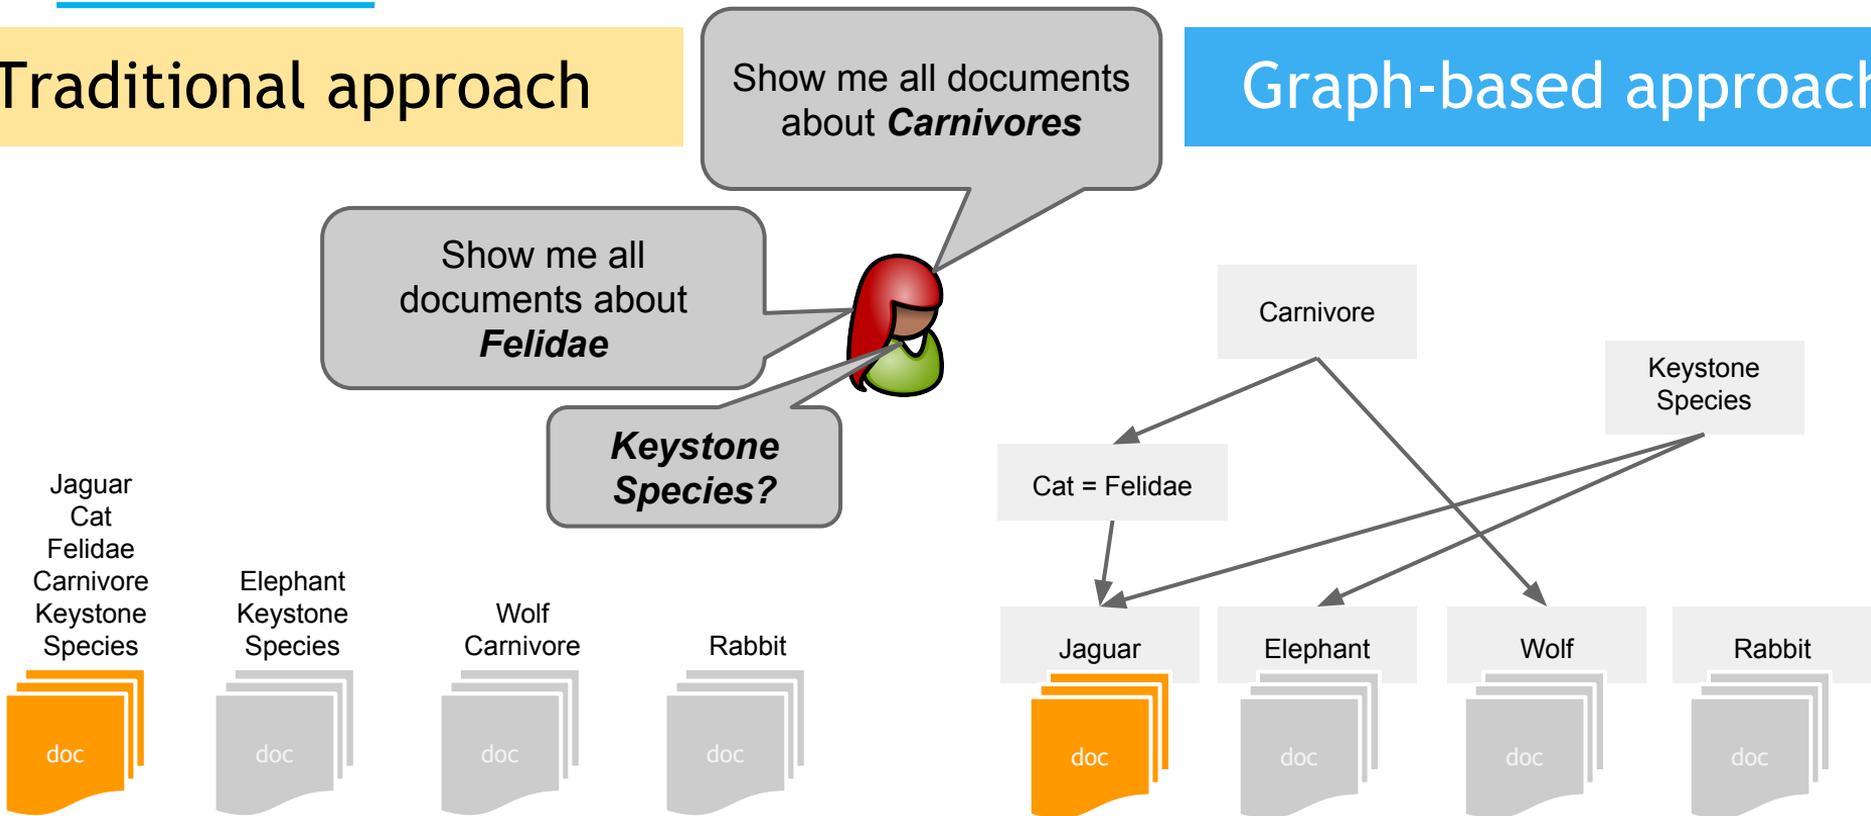


Ontological model → reasoning



Traditional approach

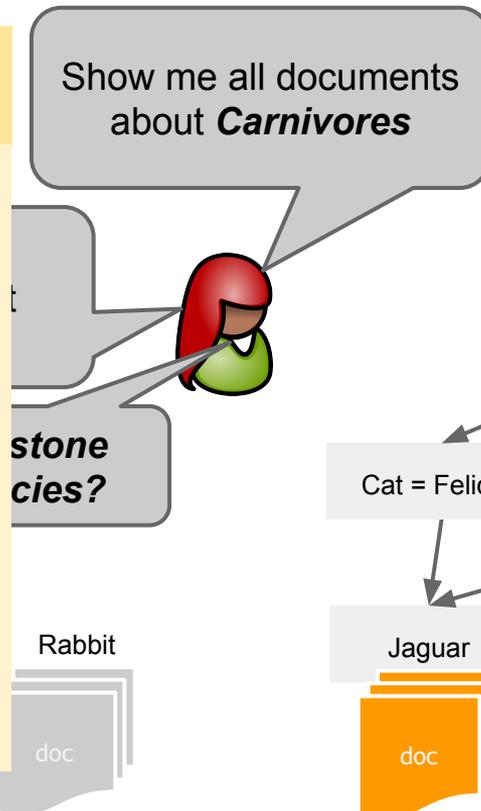
Graph-based approach



Traditional approach

Metadata per document

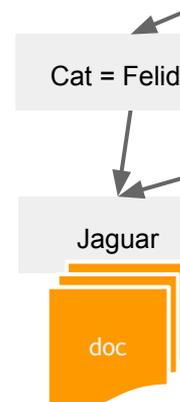
1. No or little network effects
2. No reuse of metadata
3. Metadata resides in silos
4. Data quality hard to measure
5. Not machine-readable



Graph-based approach

Knowledge about metadata

1. Explicit knowledge models
2. Reusable and measurable
3. Metadata is machine-processable
4. Standards-based metadata
5. Linkable metadata opens silos



Next steps



[PoolParty Academy](#)



[Video Tutorials](#)



[Online Help & Manuals](#)



[Demo Account](#)

Andreas Blumauer

CEO & Managing Partner, Semantic Web Company

- ▶ andreas.blumauer@semantic-web.com
- ▶ <https://www.linkedin.com/in/andreasblumauer>
- ▶ <https://twitter.com/semwebcompany>
- ▶ <https://ablvienna.wordpress.com/>

