Knowledge Graphs, Explainable Deep Learning, and What They Have To Do With Each Other

Pascal Hitzler
Data Semantics Laboratory (DaSe Lab)
Data Science and Security Cluster (DSSC)
Wright State University
http://www.pascal-hitzler.de

Data Semantics Lab

DaSe Lab

SSC

WRIGHT STATE UNIVERSITY

Some Data

- From Germany, dual citizen. PhD in Ireland (in Mathematics)
- Wright State University since 2009.
  - Assistant Professor 2009-2012
  - Associate Professor 2012-2015
  - (Full) Professor since 2015
  - Endowed NCR Distinguished Professor since 2016
- Over 400 publications
- Over 8,000 Google Scholar citations

- Previous graduate students and postdocs now at (selection):

TU Dresden, Germany	Universitas Indonesisa
UG Athens	Southeast University China
U Bonn, Germany	UN Lisboa, Portugal
UN Headquarters New York	Amazon
GE Global Research	Nuance
	etc.

Teaching Highlights

- **Redesigned discrete math sequence for computer scientists with focus on underprepared students, and increased their retention rate from 6% to 24%.**
- **Most of my classes have an additional distance learning section.**
- **I am teaching most of my classes as flipped classrooms.**
- **I received specific funding from my host institution for my teaching innovations.**

Textbook: Syntax & Semantics

DaSe Lab

Pascal Hitzler, Markus Krötzsch, Sebastian Rudolph

Foundations of Semantic Web Technologies

Chapman & Hall/CRC, 2010

Choice Magazine Outstanding Academic Title 2010 (one out of seven in Information & Computer Science)

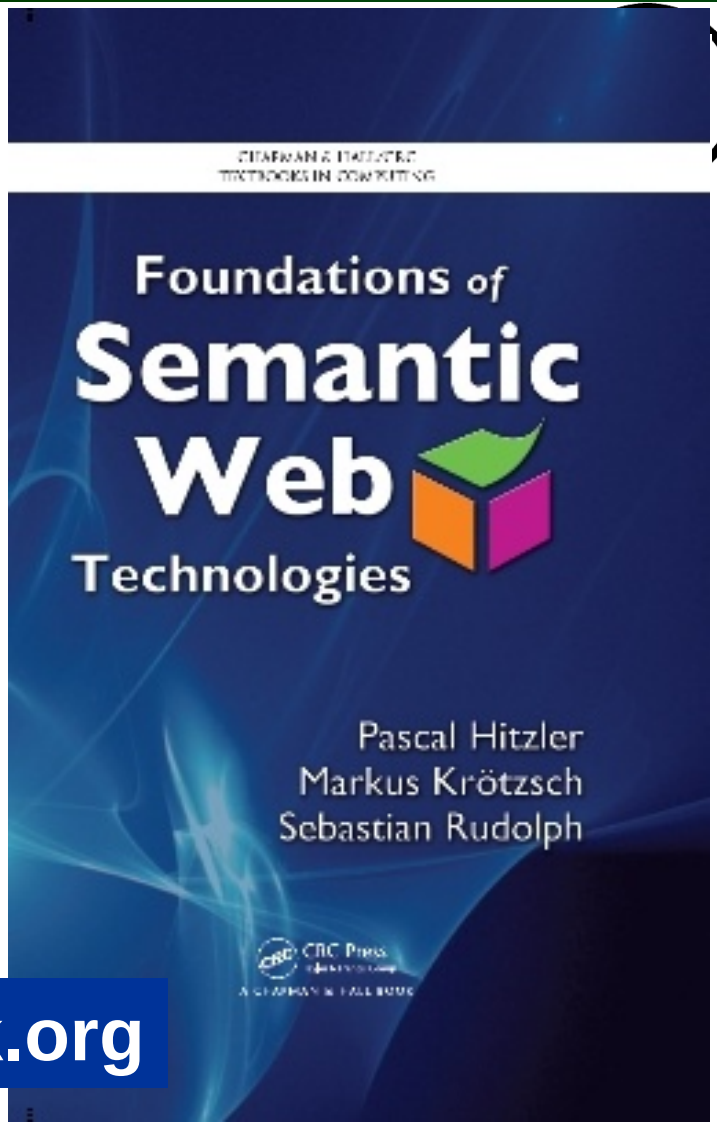

http://www.semantic-web-book.org

Semantic Web journal

- EiCs: Pascal Hitzler
  Krzysztof Janowicz
- Funded 2010
- 2017 Impact factor of 2.889, top (with 1.3 distance) of all journals with “Web” in the title

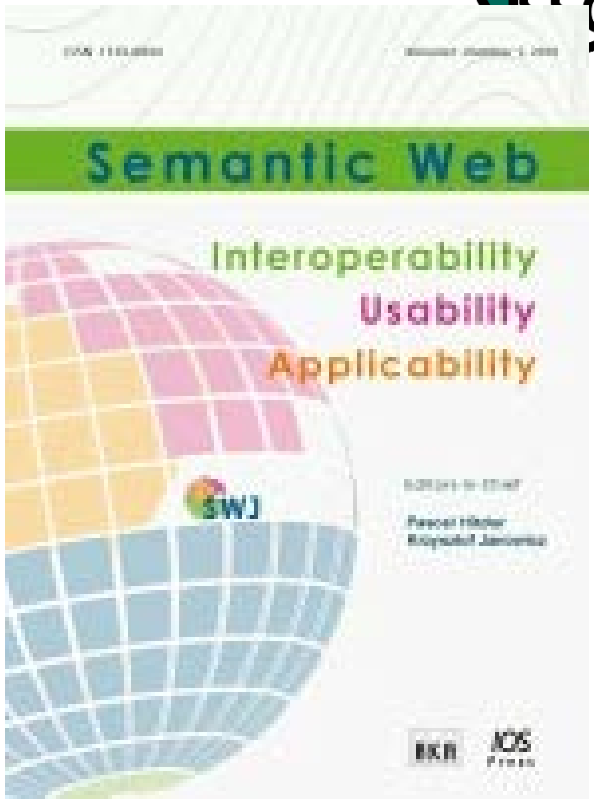

- We very much welcome contributions at the “rim” of traditional Semantic Web research – e.g., work which is strongly inspired by a different field.
- Non-standard (open & transparent) review process.

- **http://www.semantic-web-journal.net/**

Knowledge Graphs

Theresa May

British Prime Minister

tmay.co.uk

Theresa Mary May is a British politician who has served as Prime Minister of the United Kingdom and Leader of the Conservative Party since July 2016, the second woman to hold both positions. Wikipedia

Born: October 1, 1956 (age 60), Eastbourne, United Kingdom

Height: 5′ 8″

Party: Conservative Party

Spouse: Philip May (m. 1980)

Education: St Hugh's College, Oxford (1974–1977)

Previous offices: Home Secretary (2010–2016), MORE

Profiles

Twitter Facebook

People also search for View 15+

St Hugh's College, Oxford

College in Oxford, England

Website Directions

St Hugh's College is one of the constituent colleges of the University of Oxford. It is located on a 14.5-acre site on St Margaret's Road, to the north of the city centre. Wikipedia

Address: St Margaret's Rd, Oxford OX2 6LE, UK

Principal: Elish Angiolini

Phone: +44 1865 274900

Founder: Elizabeth Wordsworth

Founded: 1886

Named for: Hugh of Lincoln

Undergraduates: 432 (2011–2012)

Suggest an edit · Own this business?

Reviews from the web

4.1/5 University Rooms · 2,310 votes

Send to your phone Send

Notable alumni View 40+


Theresa May


Aung San Suu Kyi


Barbara Castle

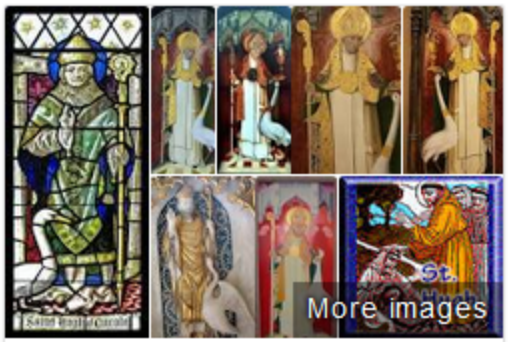

Hugh of Lincoln

Saint

Hugh of Lincoln, also known as Hugh of Avalon, was a French noble, Benedictine and Carthusian monk, bishop of Lincoln in the Kingdom of England, and Catholic saint. Wikipedia

Born: 1140, Avalon, France

Died: November 16, 1200, London, United Kingdom

Feast: 16 November (R.C.C.); 17 November (Anglican)

Major shrine: Lincoln Cathedral

Attributes: a white swan

Patronage: sick children, sick people, shoemakers and swans

People also search for


Little Saint Hugh of Lincoln


Thomas More


William Howard, 1st Visco...

NITRD action

DaSe Lab

SSC

From https://www.nitrd.gov/news/Open_Knowledge_Network.aspx:

Starting in July 2016, the Big Data Interagency Working Group (BD IWG) leadership has been involved in two meetings to discuss the viability, and possible first steps to creating a joint public/private open data network infrastructure, the Open Knowledge Network (OKN). The vision of OKN is to create an open knowledge graph of all known entities and their relationships, ranging from the macro (have there been unusual clusters of earthquakes in the US in the past six months?) to the micro (what is the best combination of chemotherapeutic drugs for a 56 y/o female with stage 3 glioblastoma and an FLT3 mutation but no symptoms of AML?). OKN is meant to be an inclusive, open, community activity resulting in a knowledge infrastructure that could facilitate and empower a host of applications and open new research avenues including how to create trustworthy knowledge networks/graphs.

The current hype

- **Often, a hype is created because something new has been established.**
- **In this case, the hype is often over before the technology has really matured to the level of application development.**

- **The current knowledge graph hype is different. Because there was already a pre-maturity hype 15 years ago, under a different name …**

Schema.org

- **Collaboratively launched in 2011 by Google, Microsoft, Yahoo, Yandex.**
  2011: 297 classes, 187 relations
  2015: 638 classes, 965 relations
- **Simple schema, request to web site providers to annotate their content with schema.org markup. Promise: They will make better searches based on this.**
- **2015: 31.3% of Web pages have schema.org markup, on average 26 assertions per page.**

Ramanathan V. Guha, Dan Brickley, Steve Macbeth: Schema.org: Evolution of Structured Data on the Web. ACM Queue 13(9): 10 (2015)

- TrainTrip
- Organization
  - Airline
  - Corporation
  - EducationalOrganization
    - CollegeOrUniversity
    - ElementarySchool
    - HighSchool
    - MiddleSchool
    - Preschool
    - School
  - GovernmentOrganization
  - LocalBusiness
    - AnimalShelter
    - AutomotiveBusiness
      - AutoBodyShop
      - AutoDealer
      - AutoPartsStore
      - AutoRental
      - AutoRepair
      - AutoWash
      - GasStation
      - MotorcycleDealer
      - MotorcycleRepair
    - ChildCare
    - Dentist
    - DryCleaningOrLaundry
    - EmergencyService
      - FireStation
      - Hospital
      - PoliceStation
    - EmploymentAgency
    - EntertainmentBusiness
      - AdultEntertainment
      - AmusementPark
      - ArtGallery
      - Casino
      - ComedyClub
      - MovieTheater
      - NightClub
    - FinancialService
      - AccountingService
      - AutomatedTeller
      - BankOrCreditUnion
      - InsuranceAgency
    - FoodEstablishment
      - Bakery
      - BarOrPub
      - Brewery
      - CafeOrCoffeeShop
      - FastFoodRestaurant

WIKIDATA

Main page
Community portal
Project chat
Create a new item
Recent changes
Random item
Query Service
Nearby
Help
Donate

Print/export
Create a book
Download as PDF
Printable version

In other projects
Wikimedia Commons
MediaWiki
Meta-Wiki
Wikispecies
Wikibooks
Wikinews
Wikipedia
Wikiquote
Wikisource
Wikiversity
Wikivoyage
Wiktionary

Tools
What links here

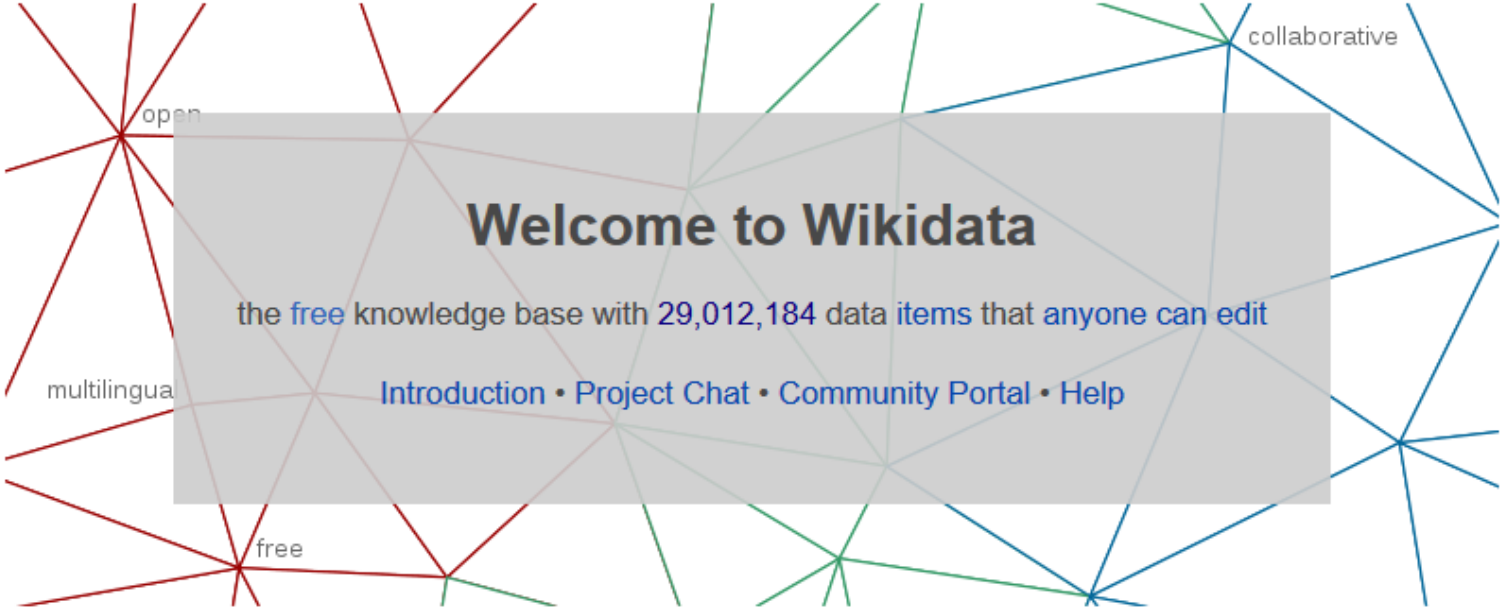

Welcome to Wikidata

the free knowledge base with 29,012,184 data items that anyone can edit

Introduction • Project Chat • Community Portal • Help

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, 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.

Linked Data

DaSe Lab

A bit older but somewhat more expressive: Linked Data on the Web

Number of Datasets	**2017-01-26**	**1,146**
in the connected	2014-08-30	570
“LOD Cloud”	2011-09-19	295
	2010-09-22	203
	2009-07-14	95
	2008-09-18	45
	2007-10-08	25
	2007-05-01	12

38.606.408.854 triples and counting!

LOD Laundromat

Linked Data: Volume

DaSe Lab

SSC

Geoindexed Linked Data – courtesy of Krzysztof Janowicz, 2012

http://stko.geog.ucsb.edu/location_linked_data

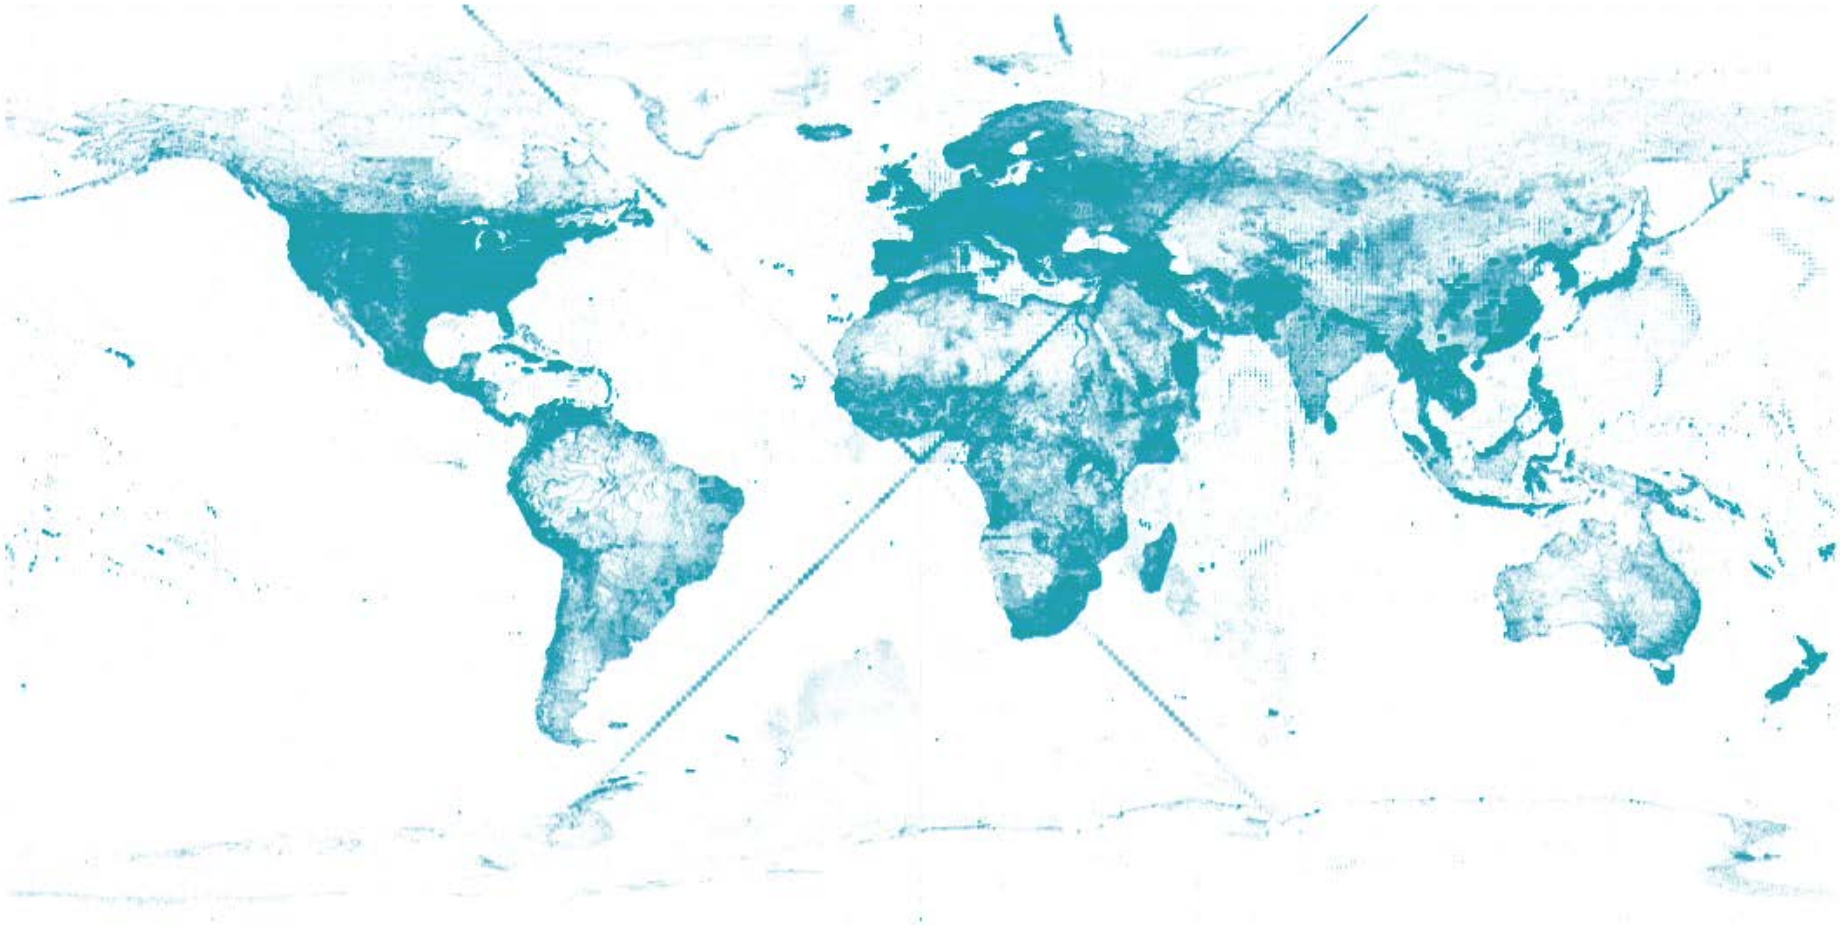

A bit of history

DaSe Lab

SSC

Before the current hype (stimulated by Google), knowledge graphs

- **have been a core artefact of study and deployment since 2007, as part of the maturing "Semantic Web Technologies" field.**
- **have been based on maturing methods and tools around the use of ontologies in the Semantic Web field, since at least 2001.**
- **have even older roots in**
  - **Artificial Intelligence, in particular related to knowledge representation and logical (deductive) reasoning**
  - **the study of terminologies (and ontologies) pre-dating the Semantic Web era**

What makes a good knowledge graph?

Principles

Goal: Easy sharing, discovery, integration, reuse

Key aspects of knowledge graphs:

- **Syntax**
- **Semantics**
- **Graph structure**
- **Tools**

Standards for syntax and semantics have been in place since at least 2004, developed by the World Wide Web Consortium (W3C).

Graph Structure

- A schema for a knowledge graph is actually also a knowledge graph, just using more abstract terms, like
  - Classes (or types) of things
    (like, Person, or Material, or Role)
  - Possible relationships between things
    (like, persons may have daughters)
  - Complex relationship assertions
    (like, every cube has 6 sides which are squares).

- A quality schema (or ontology) serves as an intermediary between data/graph structure and human conceptualization.

- A quality schema simplifies understanding and reuseability of the knowledge graph.

Ontology and Knowledge Graph

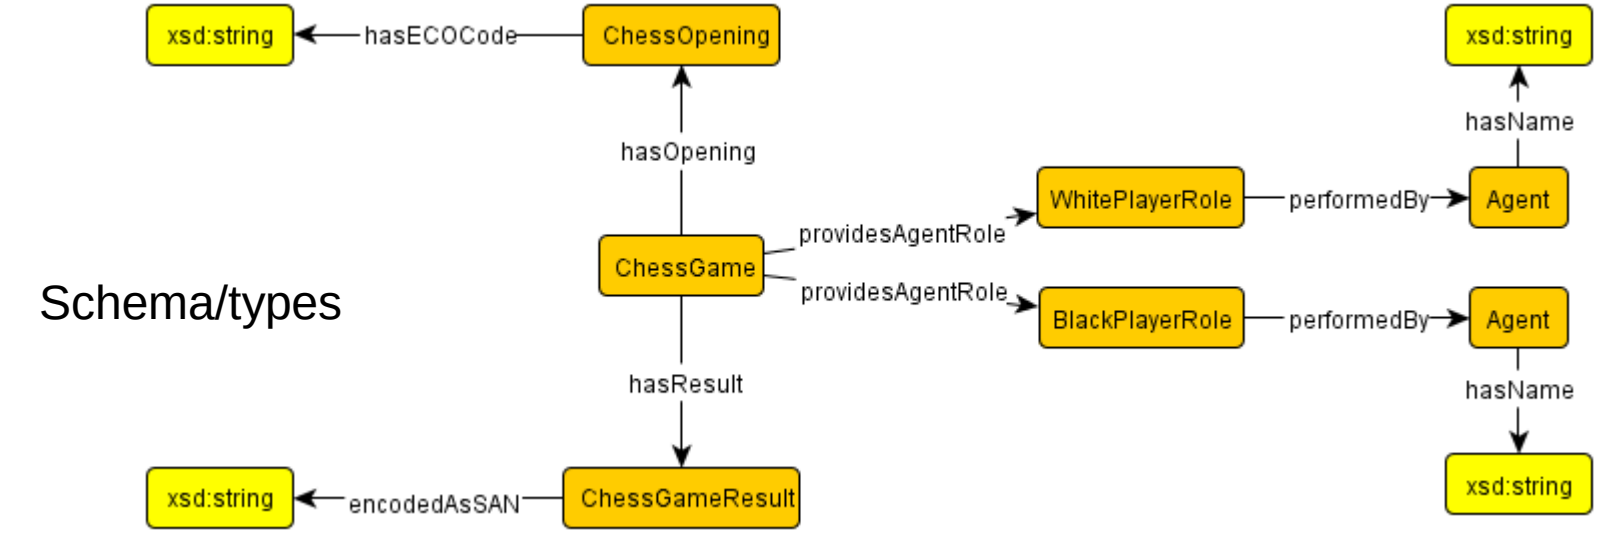

Schema/types

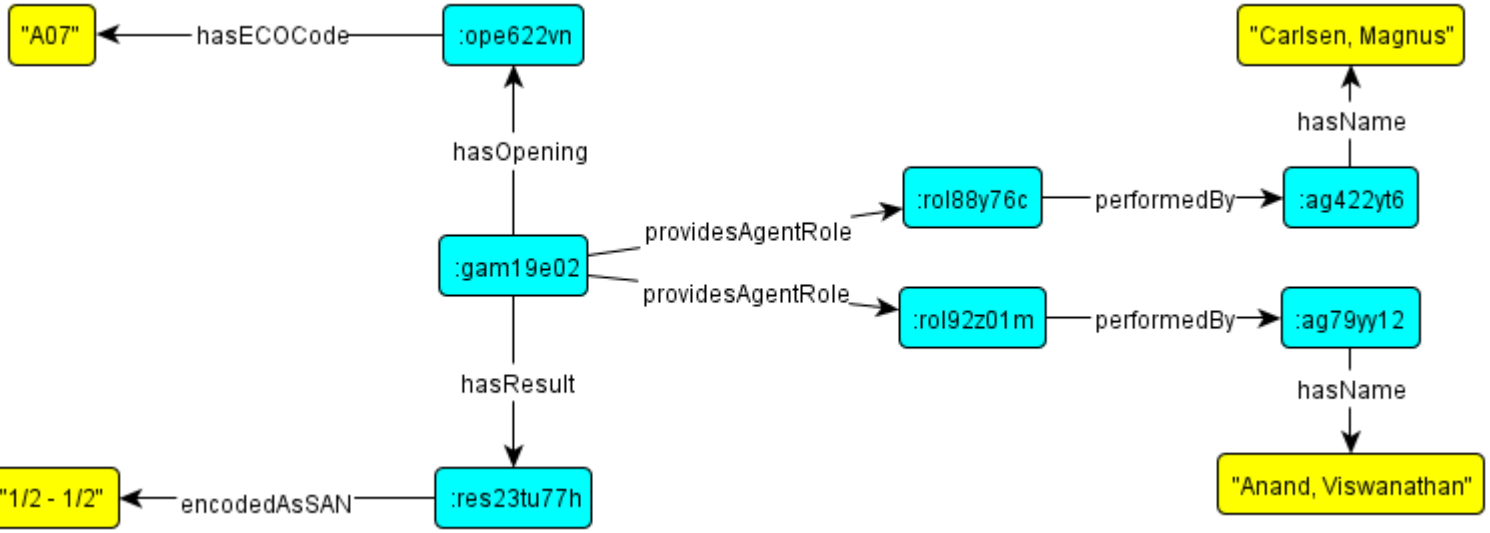

Data

Schema as a knowledge graph

DaSe Lab

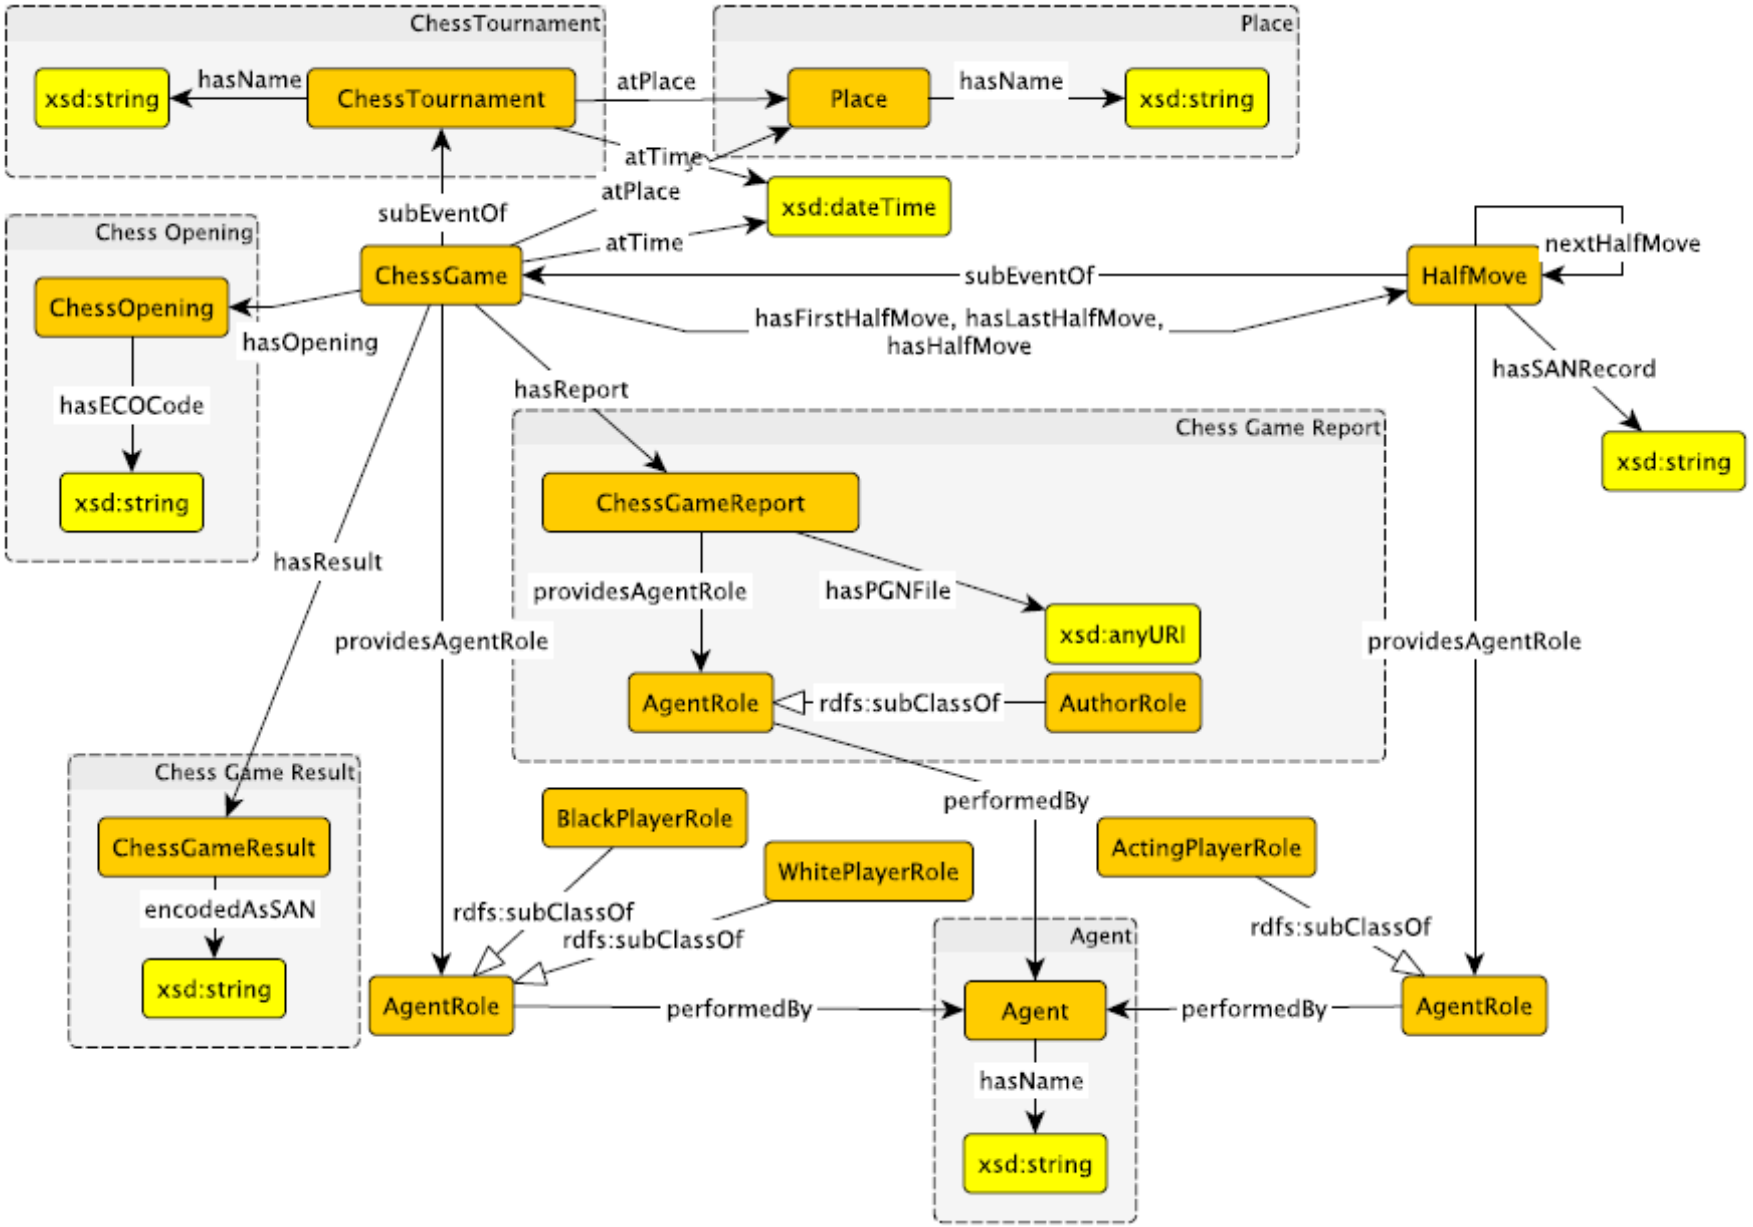

Modular Ontology Architecture

- ***Ontology Design Pattern*: A reusable ontology-piece constituting a high-quality, highly reuseable model for a commonly recurring notion.**
  E.g., "Trajectory", "Activity", "Role (of an Agent)", etc.

- **Use of well-constructed patterns minimizes risk of "naïve" modeling mistakes, thus increases reusability and repurposing of the ontology.**

- **Such ontologies are naturally made up of conceptual "modules" – these make understanding and maintenance of the ontology considerably easier.**

Schema as a knowledge graph

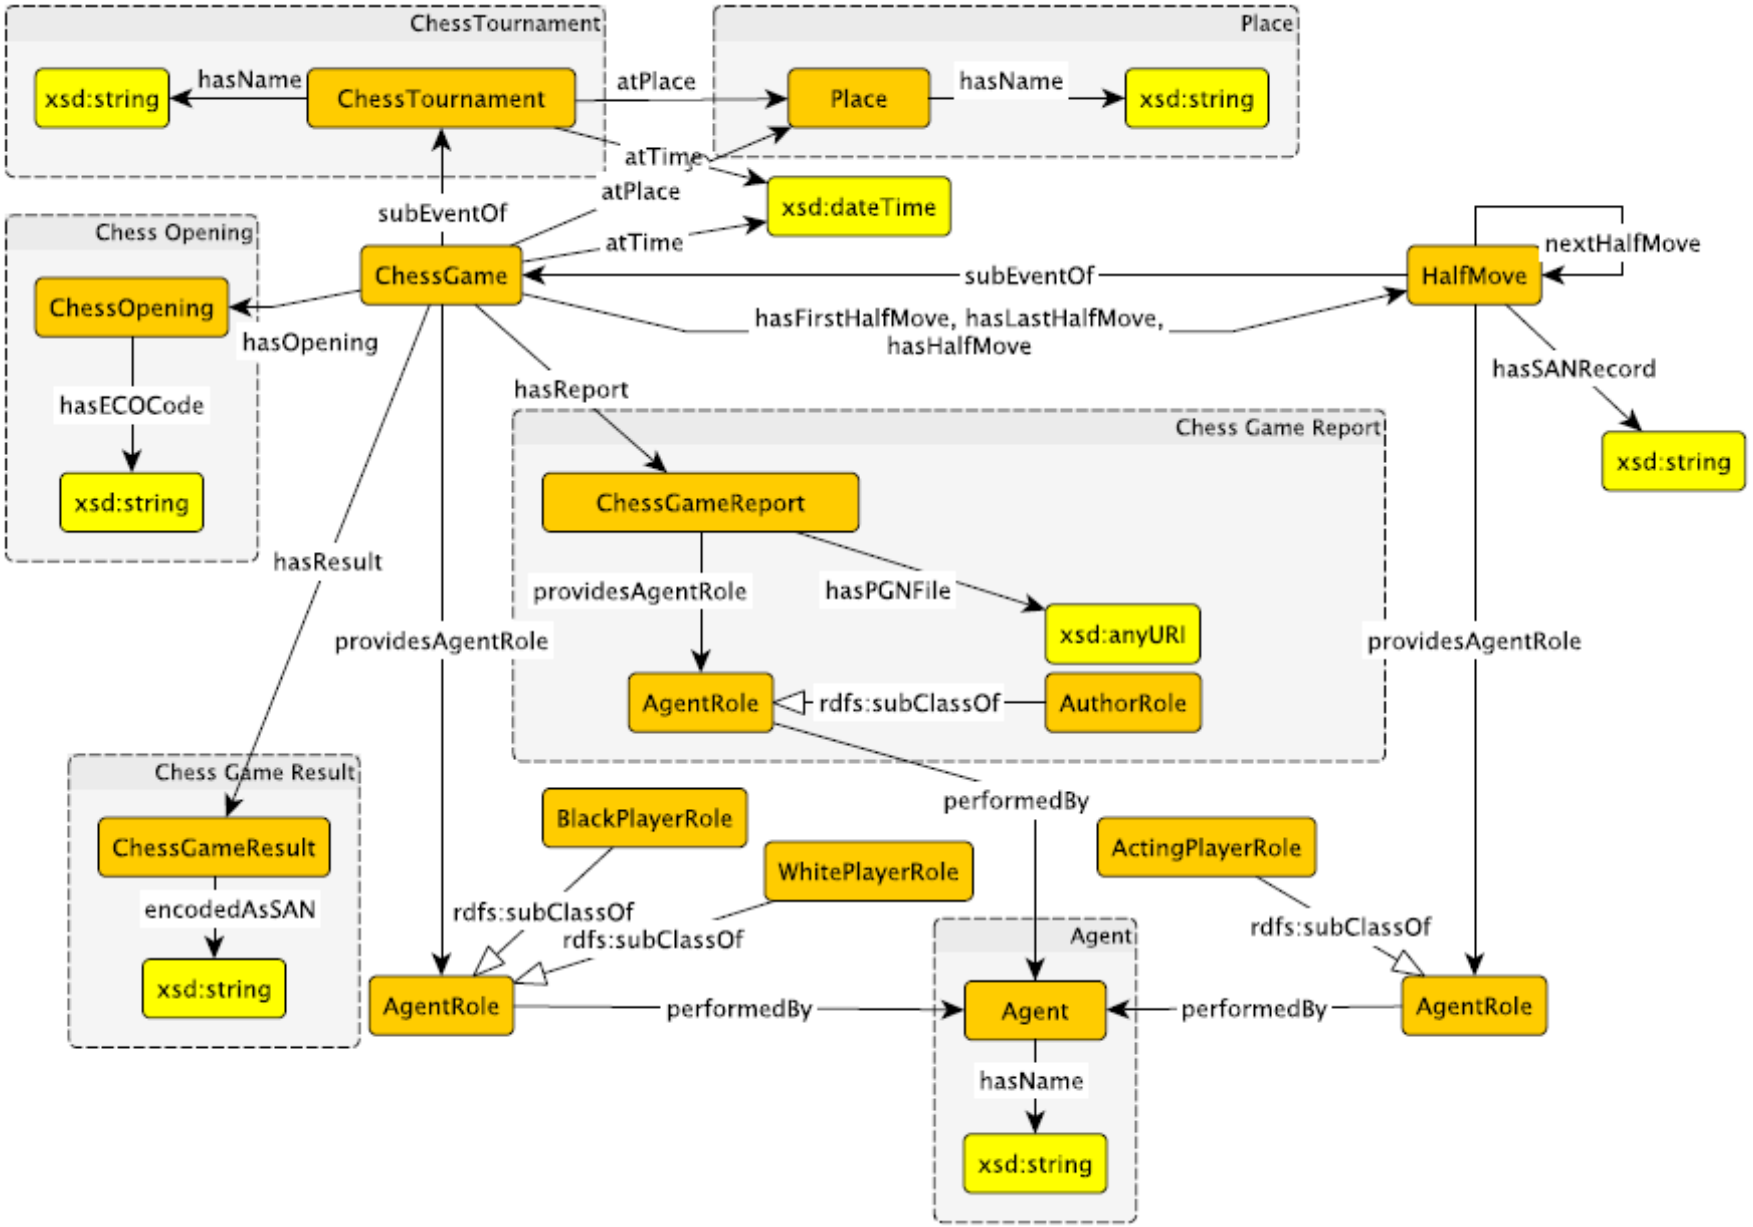

$$\mathsf{AgentRole} \sqsubseteq (=1\ \mathsf{performedBy.Agent}) \sqcap \forall \mathsf{performedBy.Agent} \tag{10.1}$$

$$\exists \mathsf{performedBy.Agent} \sqsubseteq \mathsf{AgentRole} \tag{10.2}$$

$$\top \sqsubseteq \forall \mathsf{pAR.AgentRole} \tag{10.3}$$

$$\mathsf{ChessGame} \sqsubseteq \exists \mathsf{atPlace.Place} \sqcap \forall \mathsf{atPlace.Place} \tag{10.4}$$

$$\mathsf{ChessGame} \sqsubseteq \exists \mathsf{atTime.xsd{:}dateTime} \sqcap \forall \mathsf{atTime.xsd{:}dateTime} \tag{10.5}$$

$$\mathsf{ChessGame} \sqsubseteq \exists \mathsf{pAR.BlackPlayerRole} \sqcap \exists \mathsf{pAR.WhitePlayerRole} \tag{10.6}$$

$$\exists \mathsf{subEventOf.ChessTournament} \sqcup \exists \mathsf{hasOpening.ChessOpening} \sqsubseteq \mathsf{ChessGame} \tag{10.7}$$

$$\exists \mathsf{hasResult.ChessGameResult} \sqcup \exists \mathsf{hasReport.ChessGameReport} \sqsubseteq \mathsf{ChessGame} \tag{10.8}$$

$$\mathsf{ChessGame} \sqsubseteq \forall \mathsf{subEventOf.ChessTournament} \sqcap \forall \mathsf{hasOpening.ChessOpening} \tag{10.9}$$

$$\mathsf{ChessGame} \sqsubseteq \forall \mathsf{hasResult.ChessGameResult} \sqcap \forall \mathsf{hasReport.ChessGameReport} \tag{10.10}$$

$$\mathsf{BlackPlayerRole} \sqcup \mathsf{WhitePlayerRole} \sqsubseteq \mathsf{AgentRole} \sqcap (=1\ \mathsf{pAR}^{-}\mathsf{.ChessGame}) \tag{10.11}$$

$$\mathsf{ChessGame} \sqsubseteq (=1\ \mathsf{hasFirstHalfMove.HalfMove}) \sqcap (=1\ \mathsf{hasLastHalfMove.HalfMove}) \tag{10.12}$$

$$\mathsf{ChessGame} \sqsubseteq (=1\ \mathsf{hasLastHalfMove.HalfMove}) \tag{10.13}$$

$$\mathsf{hasHalfMove} \sqsubseteq \mathsf{subEventOf}^{-} \tag{10.14}$$

$$\mathsf{hasFirstHalfMove} \sqsubseteq \mathsf{hasHalfMove} \tag{10.15}$$

$$\mathsf{hasLastHalfMove} \sqsubseteq \mathsf{hasHalfMove} \tag{10.16}$$

$$\mathsf{HalfMove} \sqsubseteq \mathsf{Event} \sqcap \exists \mathsf{pAR.ActingPlayerRole} \sqcap (=1\ \mathsf{hasHalfMove}^{-}\mathsf{.ChessGame}) \tag{10.17}$$

$$\mathsf{ActingPlayerRole} \sqsubseteq \mathsf{AgentRole} \sqcap (=1\ \mathsf{pAR}^{-}\mathsf{.HalfMove}) \tag{10.18}$$

$$\mathsf{HalfMove} \sqsubseteq (\leq 1\ \mathsf{nextHalfMove.HalfMove}) \sqcap \neg \exists \mathsf{nextHalfMove.Self} \tag{10.19}$$

$$\exists \mathsf{subEventOf.ChessGame} \sqcup \exists \mathsf{nextHalfMove.HalfMove} \sqsubseteq \mathsf{HalfMove} \tag{10.20}$$

$$\exists \mathsf{hasSANRecord.xsd{:}string} \sqsubseteq \mathsf{HalfMove} \tag{10.21}$$

$$\mathsf{HalfMove} \sqsubseteq \forall \mathsf{subEventOf.ChessGame} \sqcap \forall \mathsf{nextHalfMove.HalfMove} \tag{10.22}$$

Research Direction

- High-Quality Ontology Engineering process well understood by some experts.
- But this is “soft” knowledge. Some missing pieces:
  - Systematic exploration and evaluation of the methodology
  - Providing a powerful tool landscape supporting the methodology – plus evaluations of their effectiveness.
  - Writing it up in tutorials and textbooks, and disseminate.

- Our methods development was supported primarily through two NSF GEO projects. Several follow-on proposals are pending.

- Goal: Practical methods and tools for high-quality knowledge graph schema development.

Promise

- **Data Management (DM) is central for cost-effective / efficient data-intensive solutions, for many application areas and scenarios.**
- **DM easily takes 80% of the time when data analytics is done.**

- **Knowledge Graphs are quickly becoming a central DM tool in industry and academia.**
- **Our methods target lowering the cost of Data Management with Knowledge Graphs.**

- **I can contribute to large methods- or application-oriented projects which have Data Management components.**
- **There is also high potential for a company spin-off.**

Studies on the Semantic Web

Ontology Engineering with Ontology Design Patterns

Foundations and Applications

Pascal Hitzler, Aldo Gangemi,
Krzysztof Janowicz, Adila Krisnadhi,
Valentina Presutti (Eds.)

IOS Press

published 2016

AKA

Studies on the Semantic Web

Karl Hammar, Pascal Hitzler, Adila Krisnadhi,
Agnieszka Ławrynowicz, Andrea Giovanni Nuzzolese,
Monika Solanki (Editors)

Advances in Ontology Design and Patterns

IOS Press

published 2017

AKA

Other Aspects

Other Aspects of Knowledge Graph management we are (or have recently been) investigating:

- **Data/schema merging and integration**
- **Formal logic as schema representation language**
- **Deductive (logical) reasoning as KG engineering tool**
- **Efficient algorithms for deductive reasoning, including cloud-based**
- **KG compression**
- **Other aspects of KG quality**
- **Benchmark generation for different KG tools**

Target funding agencies

- **NSF for core new methods projects**
- **Intelligence/Defense for application-oriented projects with data management component, where effort can be used to improve and evaluate existing methods and tools.**
- **NIH similar, but haven't tapped into this yet.**
- **Potential sources also on application domains such as smart cities, data privacy and security, library science, human performance improvements, etc.**

Explaining Deep Learning

Problem

- **Connectionist systems are black boxes.**
- **How can we make sense of what they are doing? E.g. learn from what they encoded?**
  - **For human consumption.**
  - **For further processing in application software.**
- **“Explainable Artificial Intelligence” – topic currently significantly on the rise and controversially discussed.**
- **Not a new topic as such.**
  I’ve been involved in “neural-symbolic integration” research since I was a PhD student.

Understanding ANN systems

- **Explain behavior of trained (deep) NNs.**

- **Idea:**
  - **Use background knowledge in the form of linked data and ontologies to help explain.**
  - **Link inputs and outputs to background knowledge.**
  - **Use a symbolic learning system (e.g., DL-Learner) to generate an explanatory theory.**

- **We're just starting on this, I report on very first experiments.**

DL Extraction from ANNs

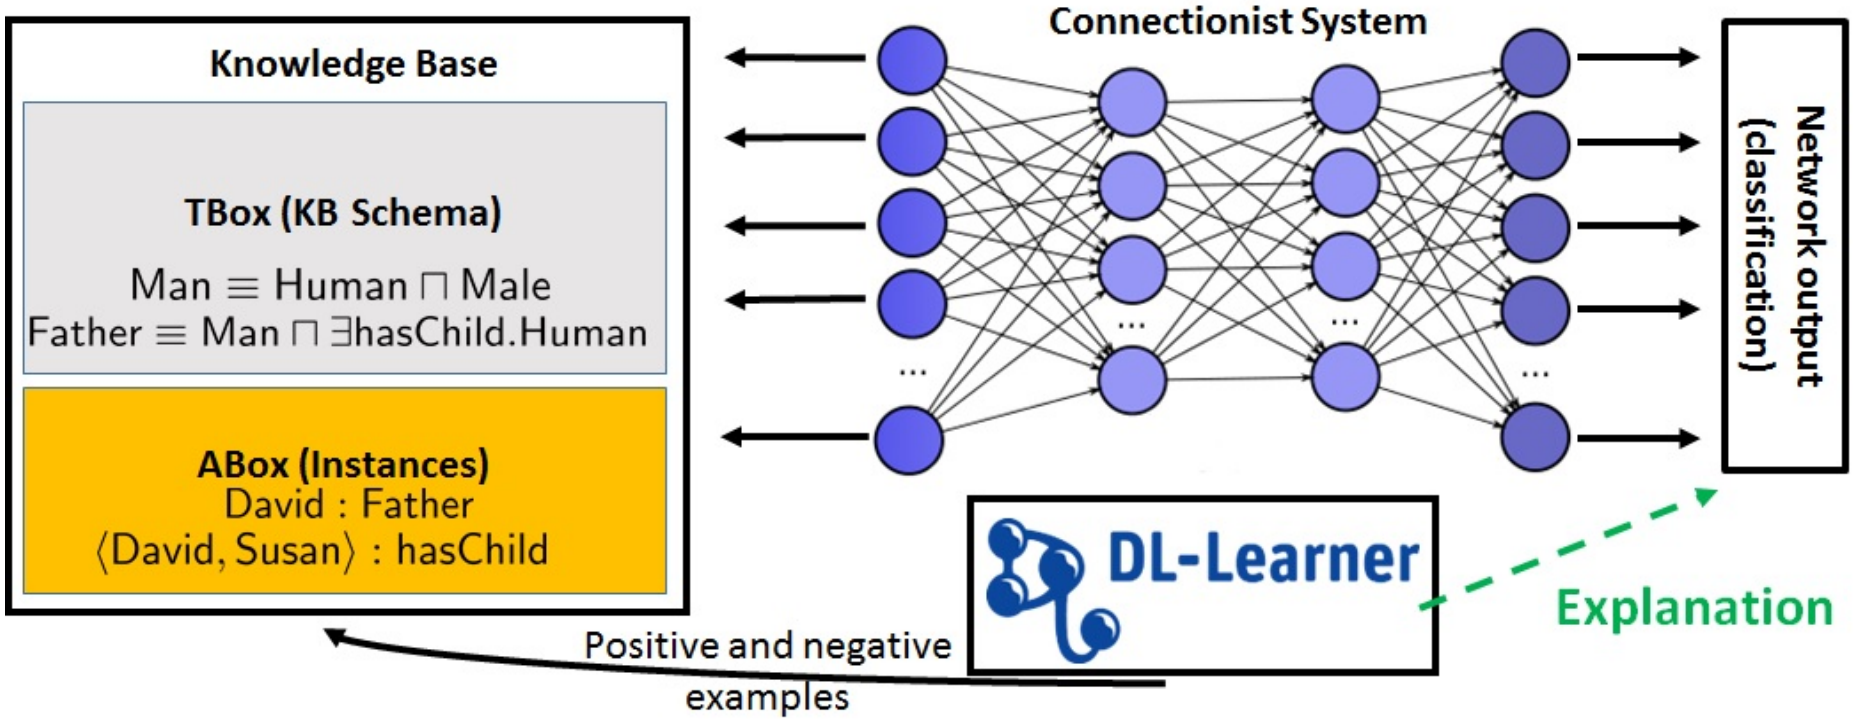

DL-Learner [Lehmann, Hitzler]

DaSe Lab

SSC

Approach similar to inductive logic programming, but using Description Logics (the logic underlying OWL).

Positive examples:

negative examples:

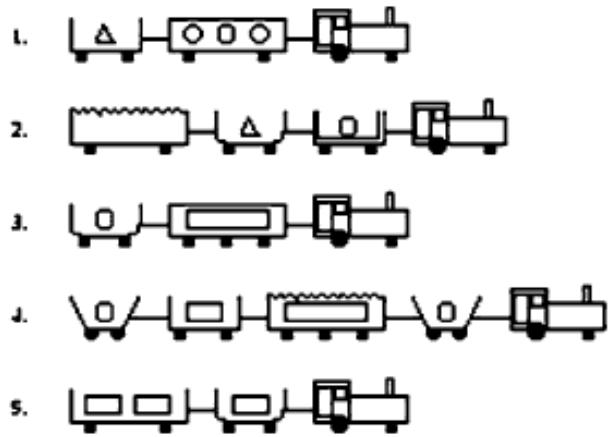

Task: find a class description (logical formula) which separates positive and negative examples.

DL-Learner

Positive examples:

negative examples:

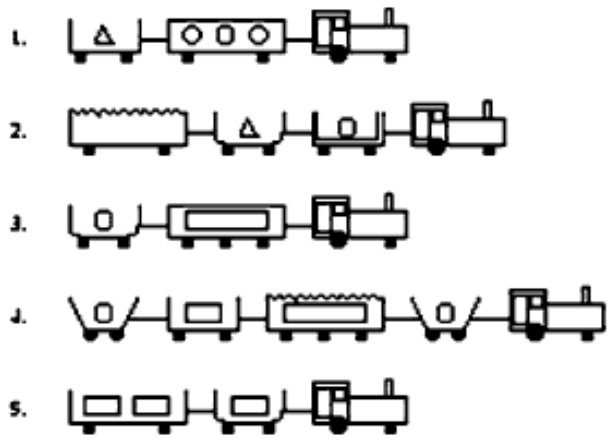

DL-Learner result: $\exists \mathtt{hasCar}.(\mathtt{Closed} \sqcap \mathtt{Short})$

In FOL:

$$\{x \mid \exists y(\mathrm{hasCar}(x, y) \wedge \mathrm{Closed}(y) \wedge \mathrm{Short}(y))\}$$

DL-Learner

DL-Learner uses refinement operators to construct ever better approximations of a solution.

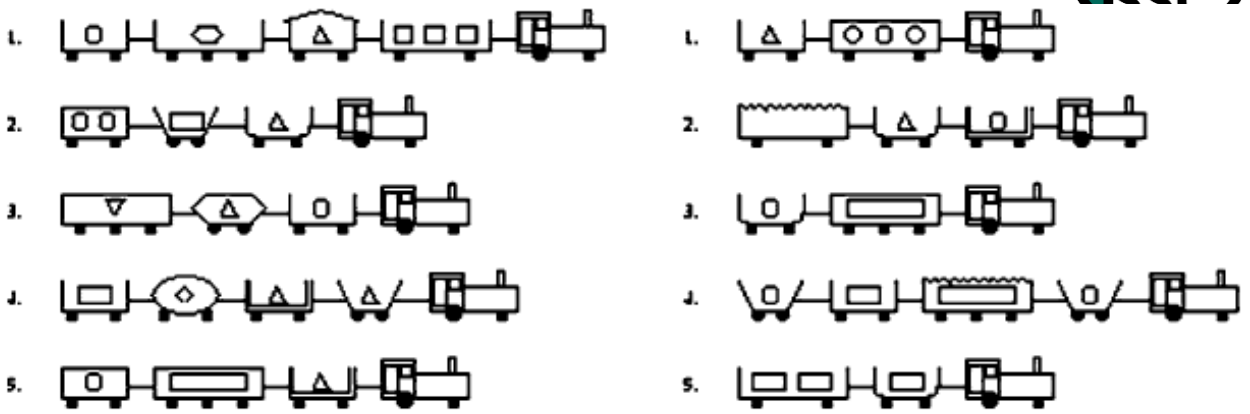

$\top$

Train – covers all examples.

$\exists \text{hasCar}.\top$

$\exists \text{hasCar}.\text{Closed}$ – covers all positives, two negatives

$\exists \text{hasCar}(\text{Closed} \sqcap \text{Short})$ – solution

Proof of Concept Experiment

DaSe Lab

SSC

Positive:

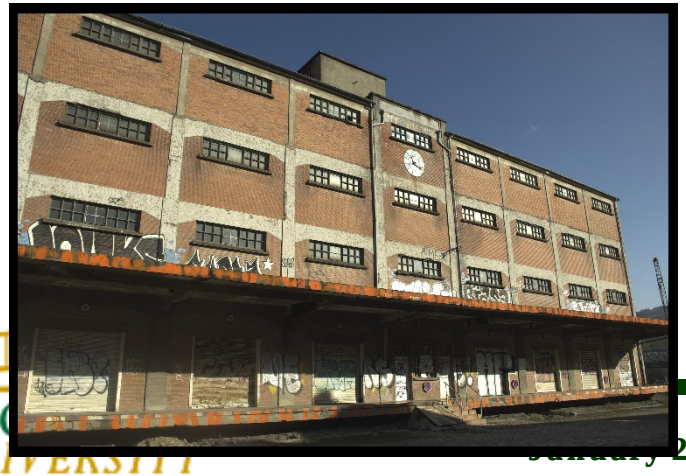

Negative:

Images

Come from the MIT ADE20k dataset
http://groups.csail.mit.edu/vision/datasets/ADE20K/
They come with annotations of objects in the picture:

```
001 # 0 # 0 # sky # sky # ""
002 # 0 # 0 # road, route # road # ""
005 # 0 # 0 # sidewalk, pavement # sidewalk # ""
006 # 0 # 0 # building, edifice # building # ""
007 # 0 # 0 # truck, motortruck # truck # ""
008 # 0 # 0 # hovel, hut, hutch, shack, shanty # hut # ""
009 # 0 # 0 # pallet # pallet # ""
011 # 0 # 0 # box # boxes # ""
001 # 1 # 0 # door # door # ""
002 # 1 # 0 # window # window # ""
009 # 1 # 0 # wheel # wheel # ""
```

Mapping to SUMO

Simple approach: for each known object in image, create an individual for the ontology which is in the appropriate SUMO class:

contains road1
contains window1
contains door1
contains wheel1
contains sidewalk1
contains truck1
contains box1
contains building1

SUMO

- **Suggested Merged Upper Ontology**
  http://www.adampease.org/OP/

- **Approx. 25,000 common terms covering a wide range of domains**

- **Centrally, a relatively naïve class hierarchy.**

- **Objects in image annotations became individuals (constants), which were then typed using SUMO classes.**

DL-Learner input

Positive:

- **img1:** road, window, door, wheel, sidewalk, truck, box, building
- **img2:** tree, road, window, timber, building, lumber
- **img3:** hand, sidewalk, clock, steps, door, face, building, window, road

Negative:

- **img4:** shelf, ceiling, floor
- **img5:** box, floor, wall, ceiling, product
- **img6:** ceiling, wall, shelf, floor, product

DL-Learner results include:

$$\exists \text{contains.Transitway}$$

$$\exists \text{contains.LandArea}$$

Proof of Concept Experiment

DaSe Lab

Positive:

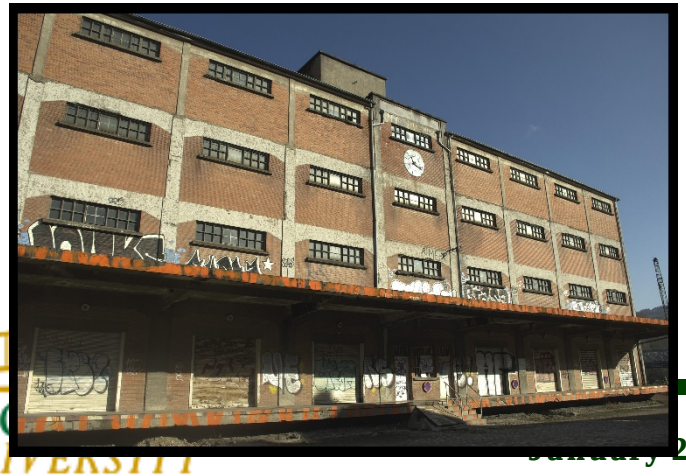

Negative:

$\exists contains.Transitway$

$\exists contains.LandArea$

First 10 DL-Learner responses

$\exists$contains.Window	(1)	$\exists$contains.LandTransitway	(6)
$\exists$contains.Transitway	(2)	$\exists$contains.LandArea	(7)
$\exists$contains.SelfConnectedObject	(3)	$\exists$contains.Building	(8)
$\exists$contains.Roadway	(4)	$\forall$contains.$\neg$Floor	(9)
$\exists$contains.Road	(5)	$\forall$contains.$\neg$Ceiling	(10)

Experiment 2

DaSe Lab

SSC

Positive (selection):

Negative (selection):

$$\exists \text{contains}.(\text{DurableGood} \sqcap \neg \text{ForestProduct})$$

Experiment 3

DaSe Lab

Positive:

Negative:

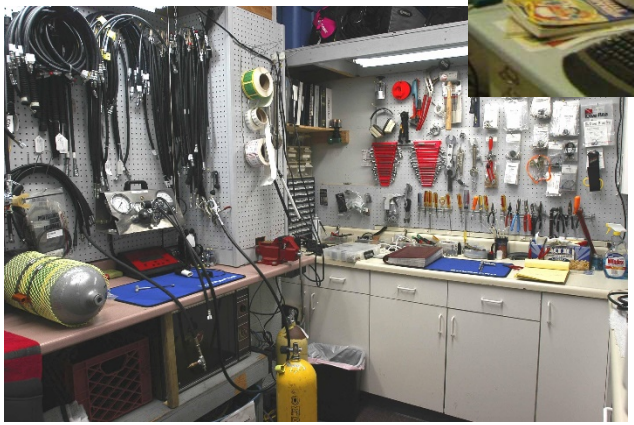

$$\forall \text{contains}.(\neg \text{Furniture} \sqcap \neg \text{IndustrialSupply})$$

Experiment 4

DaSe Lab

SSC

Positive (selection):

Negative (selection):

$\exists$contains.SentientAgent

Experiment 5

Positive:




Negative (selection):



$\exists contains.BodyOfWater$

Next steps

- Utilize more sophisticated ontology.
- Utilize more sophisticated mappings.

- Explain hidden neurons.
- Then also tune learning algorithms towards easier interpretability.

- Tune DL-Learner better to the specific task.

- Dig deeper on the ADE20k dataset case, or transfer to other datasets of media. E.g. collaboration with IBM has already been established to look into drug-drug-interaction data.

Uniqueness of approach

- **To the best of our knowledge, nobody else is pursuing explainable deep learning though background knowledge.**

- **To the best of our knowledge, nobody funded under the DARPA XAI program is pursuing explainable deep learning through background knowledge.**

Other Aspects

Other Aspects of Deep Learning we are investigating include:

- **Deep Learning methods for data integration**
- **Deep Learning for text analysis**
- **Deep Learning algorithms to support KG engineering tools (e.g., deductive reasoning and graph completion)**

- **Explaining other statistical approaches (not only deep learning) by transferring our methods.**

Target funding agencies

- **NSF for core new methods projects**
- **Intelligence/Defense for application-oriented projects regarding the use of explainable deep learning.**
- **NIH similar, but haven't tapped into this yet. We just started a collaboration with IBM TJ Watson on applying our method to drug-drug-interaction.**
- **Potential sources also on any current or emerging application domain of deep learning, including security, social media analysis, intelligence analysis, etc.**

Summary

- My work has many facets.
- My work is in synch with several current trends, including
  - Knowledge Graphs
  - Deep Learning
  - Big Data
  - Data Science
- Covering methods/foundations and applications; and the transfer between them.

- Broad options for obtaining research funding.
- Easier because of already significant visibility and standing. Since I became a US citizen (summer 2017) I also made significant inroads for defense funding, in particular establishing a network of contacts.

Thanks!