UNIVERSITÄT MANNHEIM

Knowledge Graphs
Public Knowledge Graphs

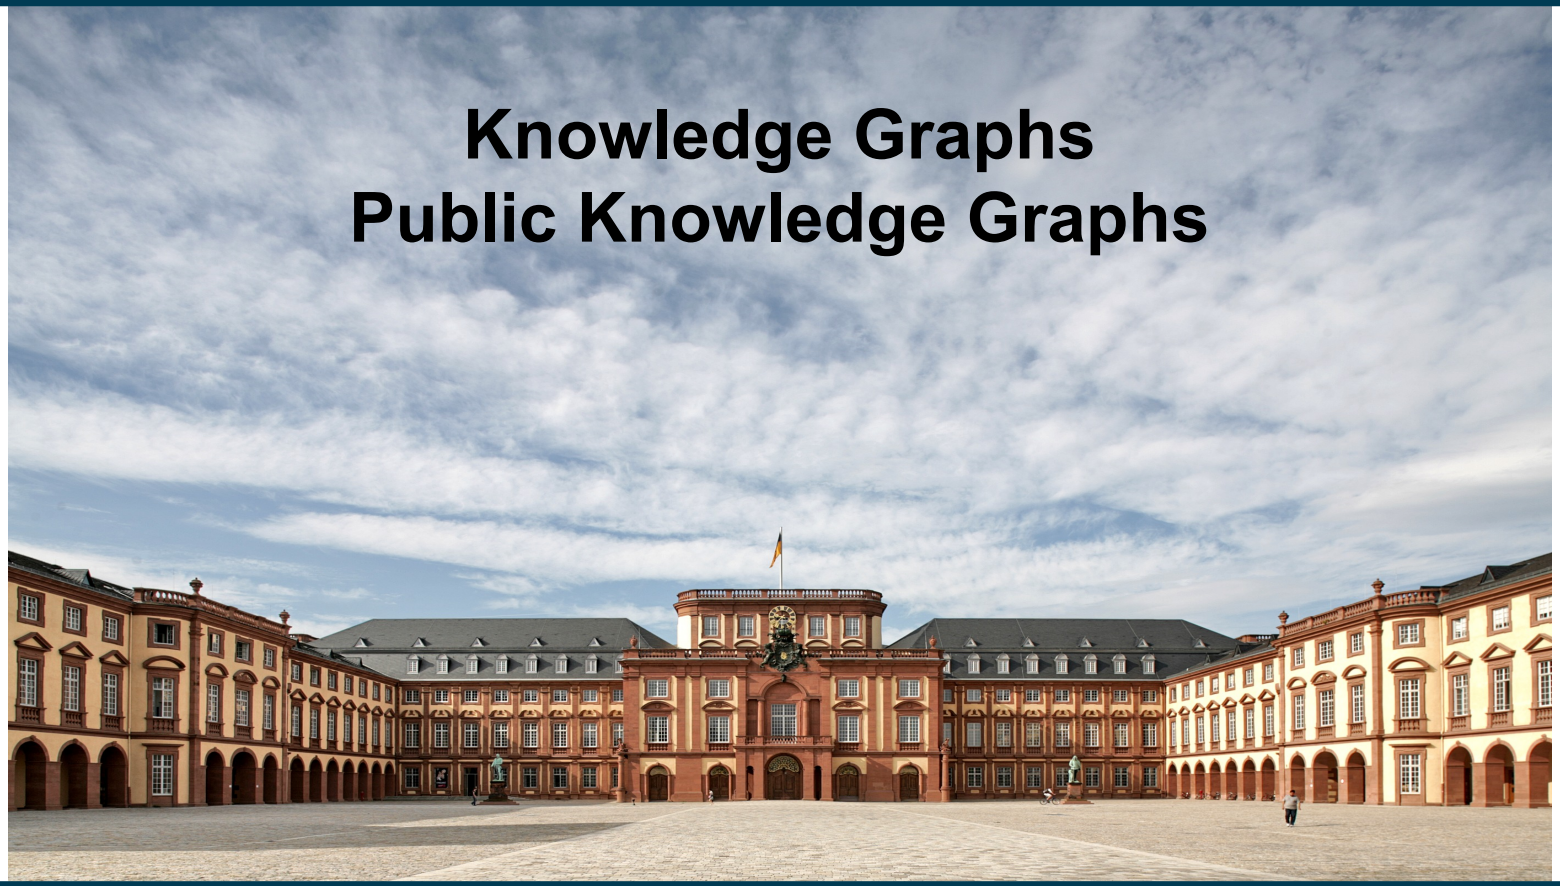

Heiko Paulheim

Previously on “Knowledge Graphs”

- Principles:
  - RDF, RDF-S, SPARQL & co
  - Linked Open Data

- Today
  - A closer look on actually existing knowledge graphs
  - Some useful, large-scale resources

Introduction

- Knowledge Graphs out there (not guaranteed to be complete)

Name	Instances	Facts	Types	Relations
DBpedia (English)	4,806,150	176,043,129	735	2,813
YAGO	4,595,906	25,946,870	488,469	77
Freebase	49,947,845	3,041,722,635	26,507	37,781
Wikidata	15,602,060	65,993,797	23,157	1,673
NELL	2,006,896	432,845	285	425
OpenCyc	118,499	2,413,894	45,153	18,526
Google's Knowledge Graph	570,000,000	18,000,000,000	1,500	35,000
Google's Knowledge Vault	45,000,000	271,000,000	1,100	4,469
Yahoo! Knowledge Graph	3,443,743	1,391,054,990	250	800

public

private

Paulheim: *Knowledge graph refinement: A survey of approaches and evaluation methods.* Semantic Web 8:3 (2017), pp. 489-508

Knowledge Graph Creation: CyC

- The beginning
  - Encyclopedic collection of knowledge
  - Started by Douglas Lenat in 1984
  - Estimation: 350 person years and 250,000 rules should do the job
    of collecting the essence of the world's knowledge

- The present (as of June 2017)
  - ~1,000 person years, $120M total development cost
  - 21M axioms and rules
  - Used to exist until 2017

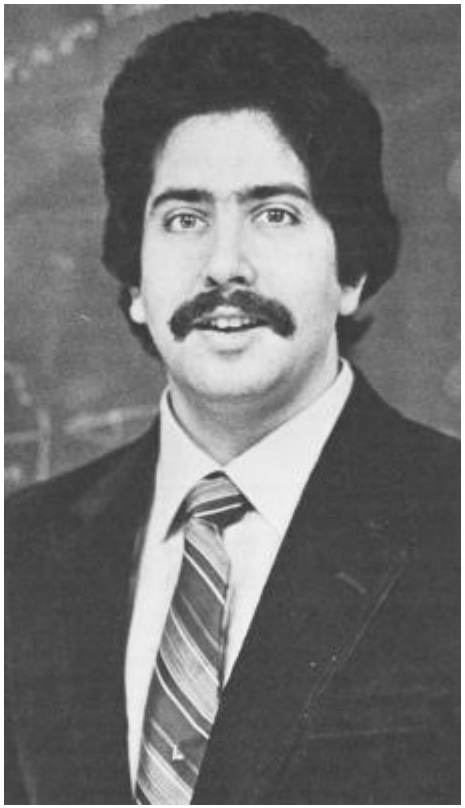

Knowledge Graph Creation: CyC

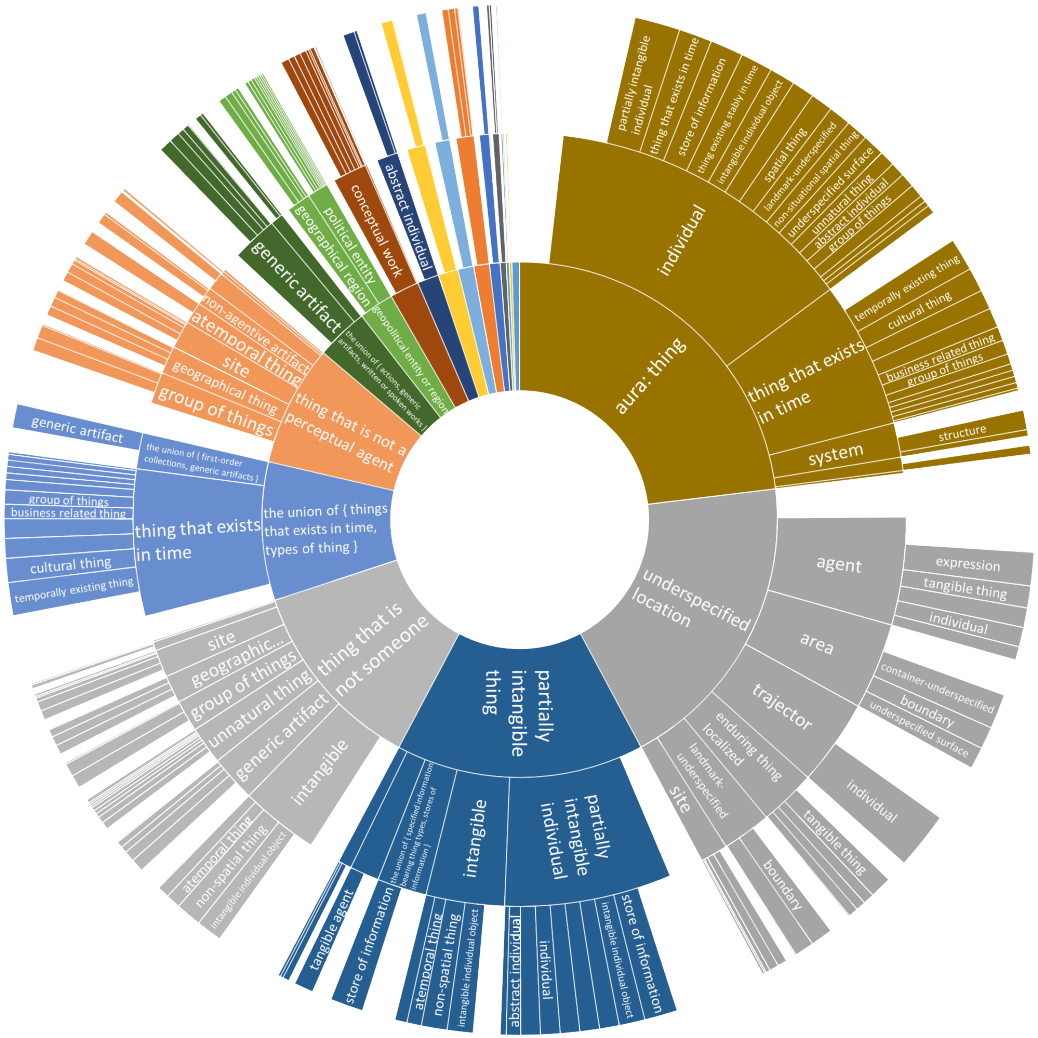

Knowledge Graph Creation

- Lesson learned no. 1:
  - Trading efforts against accuracy

Knowledge Graph Creation: Freebase

- The 2000s
  - Freebase: collaborative editing
  - Schema not fixed

coming up soon:
was it a good deal or not?

- Present
  - Acquired by Google in 2010
  - Powered first version of Google's Knowledge Graph
  - Shut down in 2016
  - Partly lives on in Wikidata (see in a minute)

Knowledge Graph Creation: Freebase

- Community based
- Like Wikipedia, but more structured

Arnold Schwarzenegger

Discuss "Arnold Schwarzenegger" Show Empty Fields

image 1 of 1

Types: Person (People), US Politician (Government), Film actor (Film), Film producer (Film), Pro Athlete (Sports), Sports Award Winner (Sports)

Also known as: Arnold Alois Schwartzenegger, The Governator

Gender: Male

Date of Birth: Jul 30, 1947

Place of Birth: Thal, Austria

Country Of Nationality: United States

Profession: Politician, Bodybuilder, Entrepreneur, Actor

Religion: Roman Catholicism

Parents: Aurelia Jadrny Schwarzenegger, Gustav Schwarzenegger

Children: Christopher Schwarzenegger, Patrick Schwarzenegger, Christina Schwarzenegger, Katherine Schwarzenegger

Siblings: Meinhard Schwarzenegger

Spouse (or domestic partner): Maria Shriver • Apr 26, 1986

Height: 1.88 m

IMDB Entry: http://www.imdb.com/name/nm0000216/

Career Start: 1968

Career End: 1980

Knowledge Graph Creation

- Lesson learned no. 2:
  - Trading formality against number of users

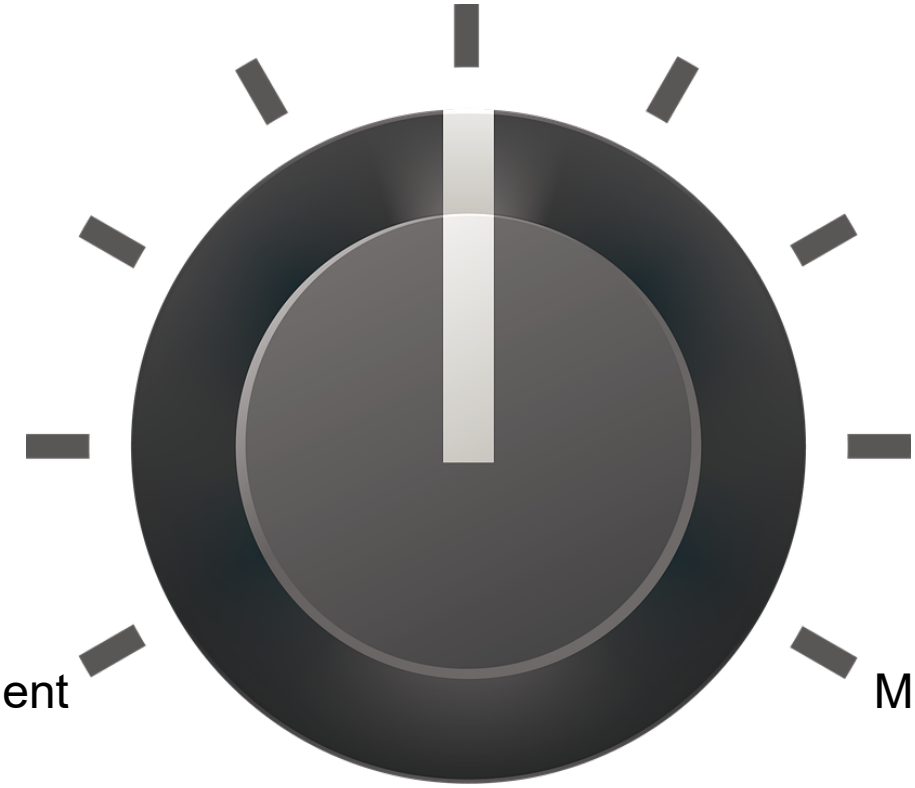

Max. user involvement

Max. degree of formality

Knowledge Graph Creation: Wikidata

- The 2010s
  - Wikidata: launched 2012
  - Goal: centralize data from Wikipedia languages
  - Collaborative
  - Imports other datasets

- Present
  - One of the largest public knowledge graphs (see later)
  - Includes rich provenance

Knowledge Graph Creation: Wikidata

- Collaborative editing

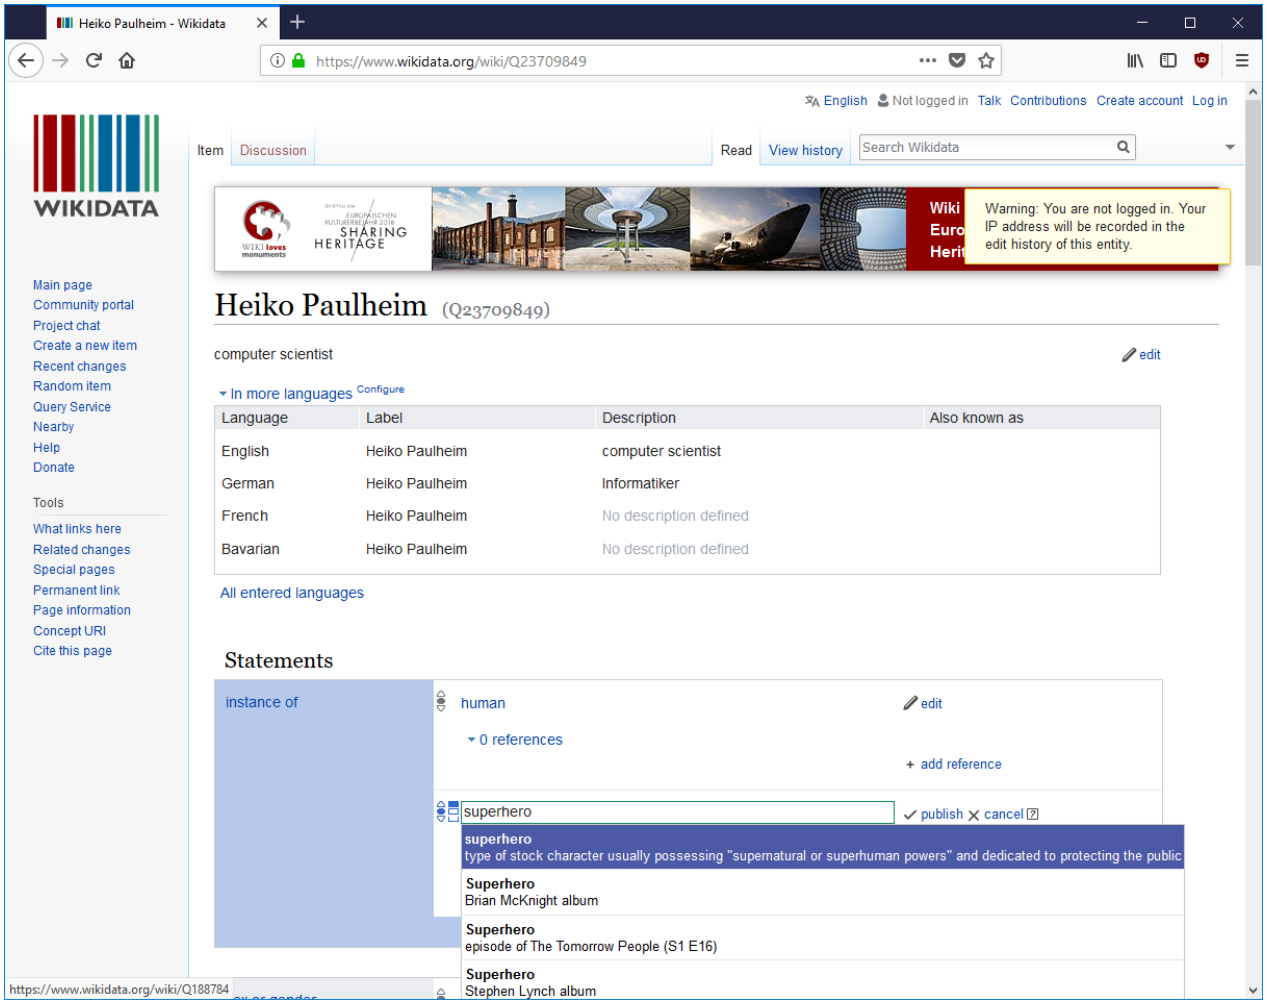

Knowledge Graph Creation: Wikidata

- Provenance

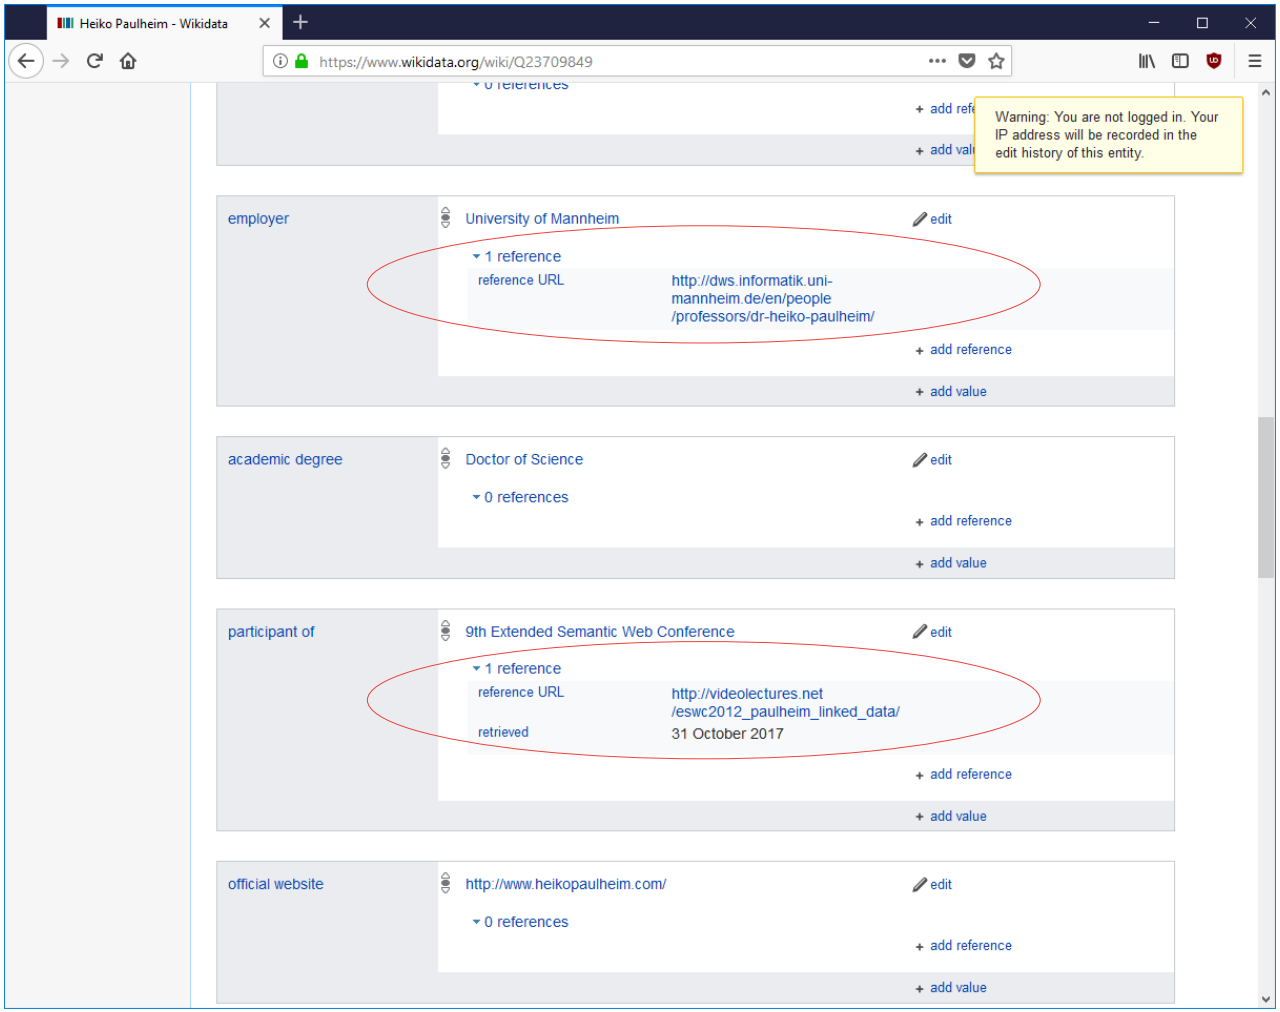

Wikidata

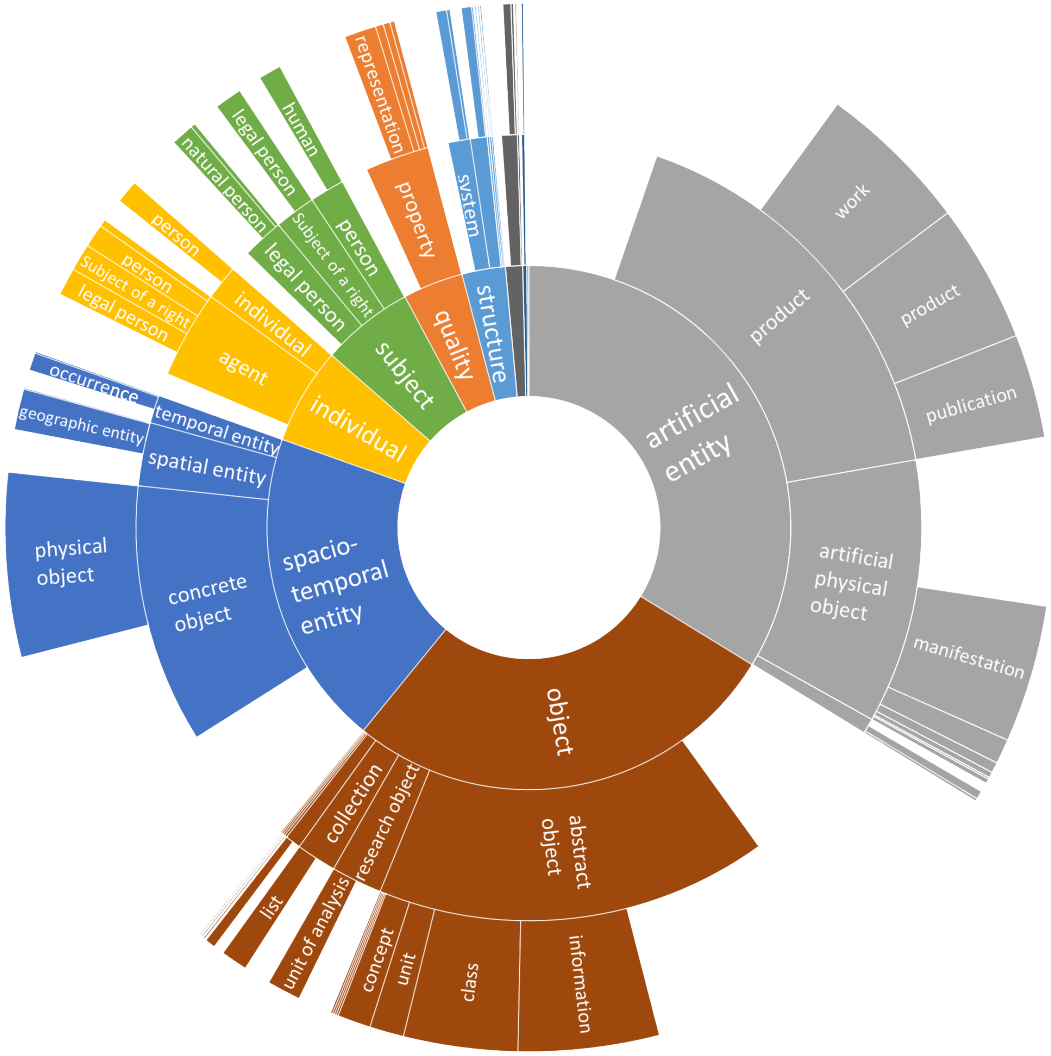
artificial entity
product
work
product
publication
artificial physical object
manifestation
object
abstract object
information
class
unit
concept
research object
collection
unit of analysis
list
spacio-temporal entity
concrete object
physical object
spatial entity
temporal entity
geographic entity
occurrence
individual
agent
person
person
Subject of a right
legal person
subject
legal person
Subject of a right
person
natural person
legal person
human
quality
property
representation
structure
system

Knowledge Graph Creation

- Lesson learned no. 3:
  - There is not one truth (but allowing for plurality adds complexity)

Knowledge Graph Creation: DBpedia & YAGO

- The 2010s
  - DBpedia: launched 2007
  - YAGO: launched 2008
  - Extraction from Wikipedia using mappings & heuristics

- Present
  - Two of the most used knowledge graphs
  - Next to Wikidata

yago
select knowledge

DBpedia

Wikipedia as a Knowledge Graph

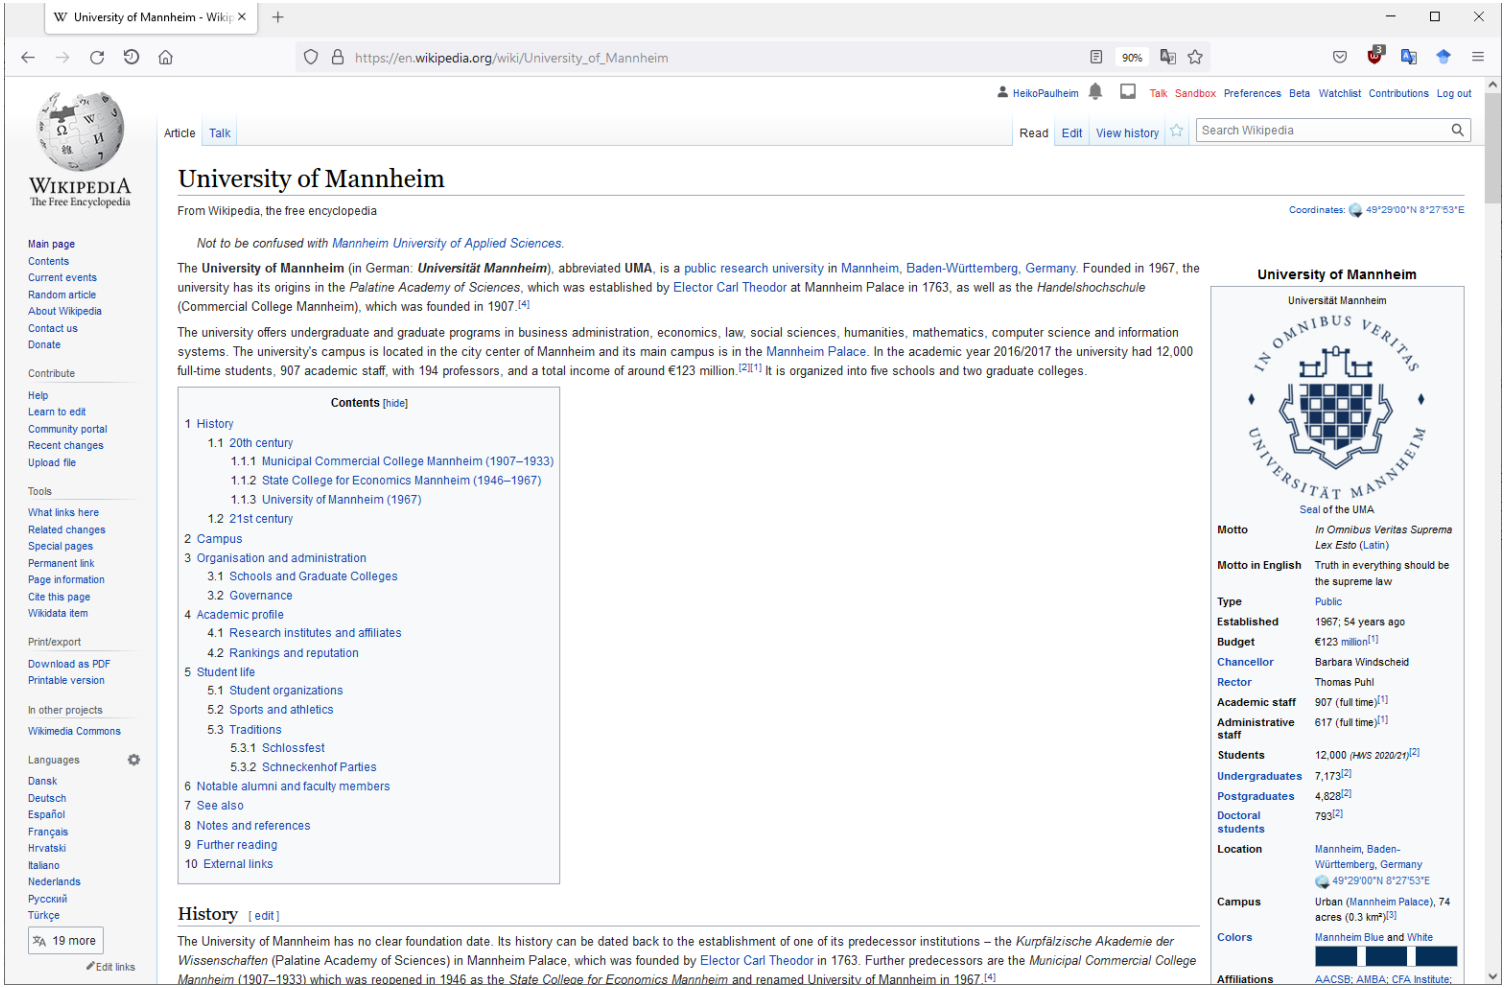

Wikipedia as a Knowledge Graph

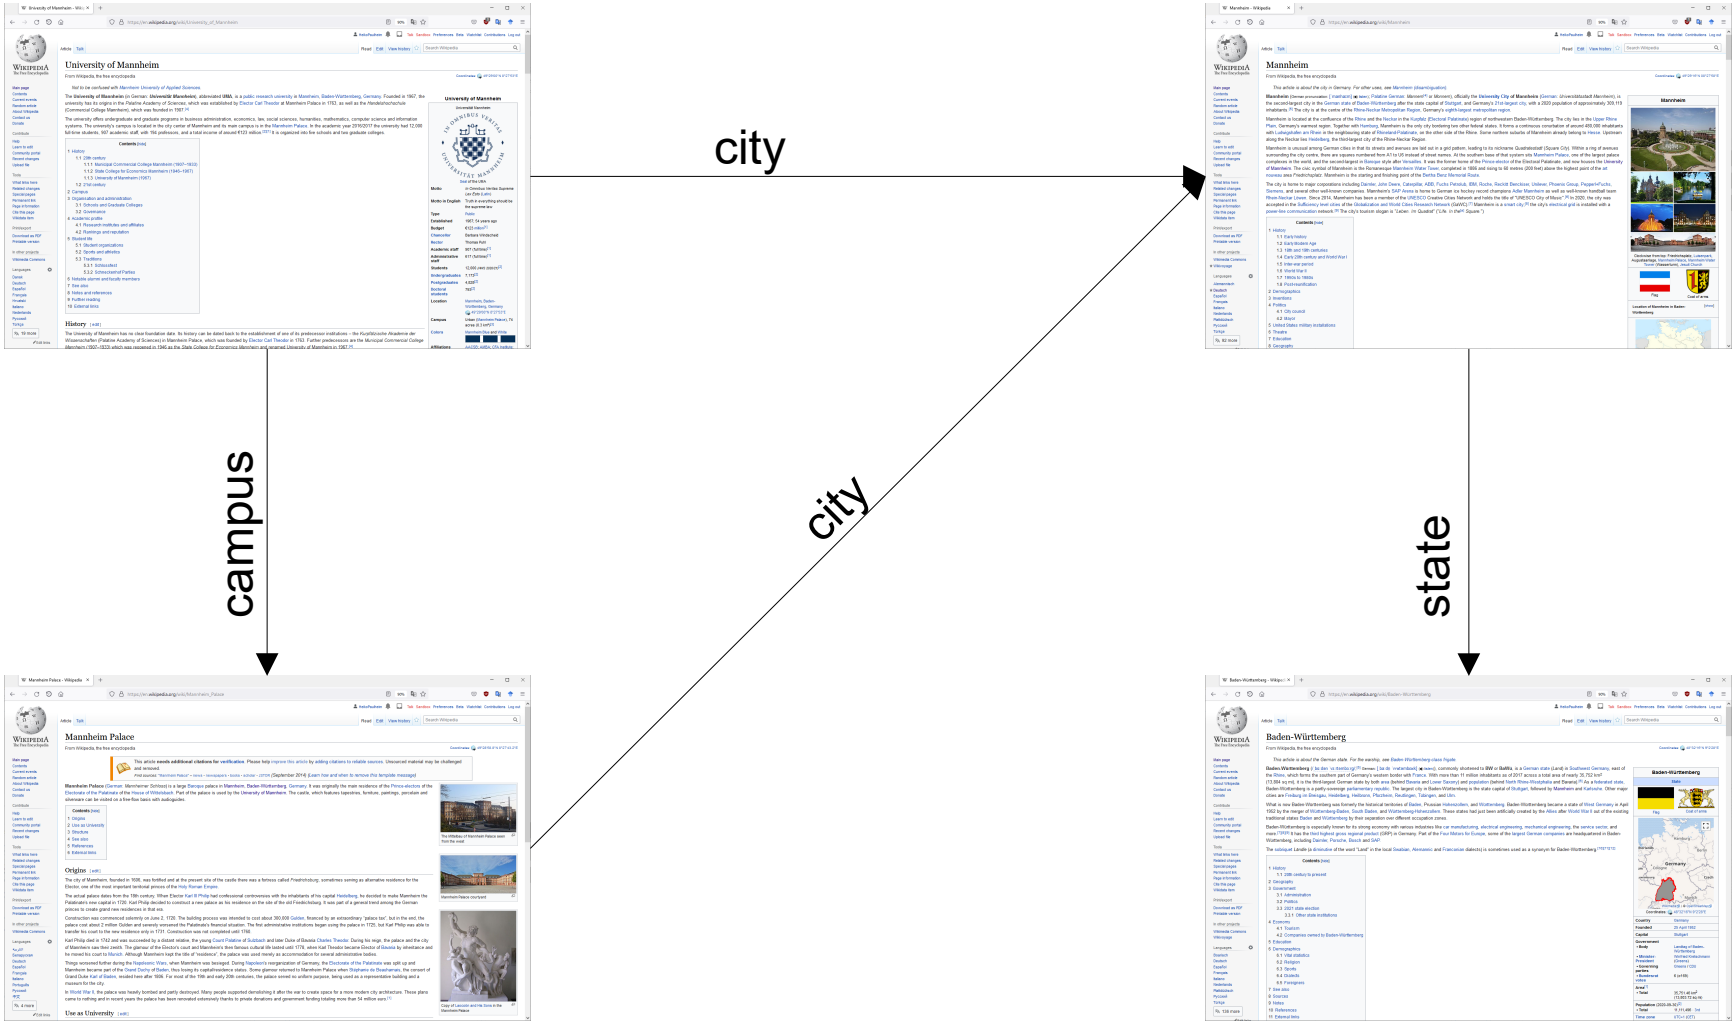

Wikipedia as a Knowledge Graph

- Mapping to a central schema/ontology

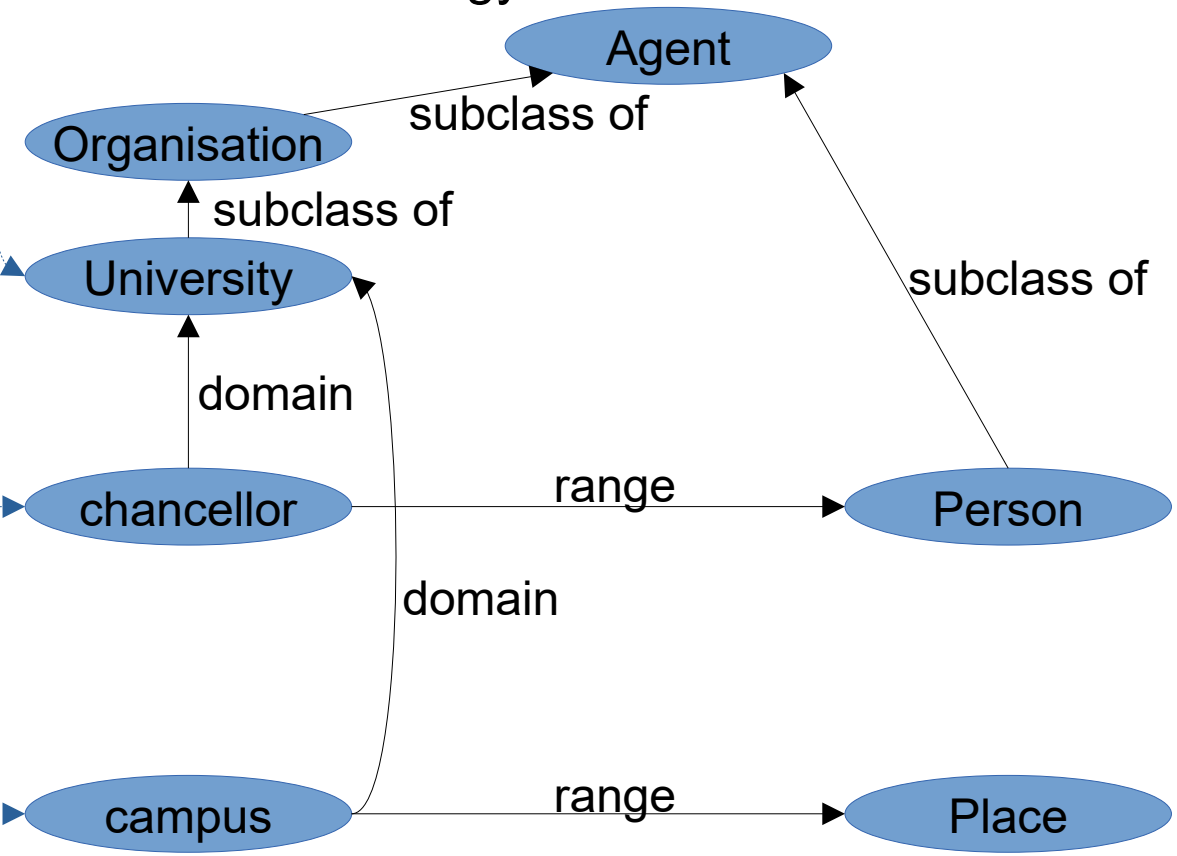

DBpedia

University of Mannheim

Universität Mannheim

Motto	
Motto in Englis	
Established	
Type	
Endowment	
Chancellor	
Rector	
Academic staf	
Administrative staff	
Students	
Undergraduates	6,915[1]
Postgraduates	4,965[1]
Doctoral students	249[1]

```
{{Infobox university
|motto            =''In Omnibus Veritas Suprema Lex Esto'' ([[Latin]])
|mottoeng         = Truth in everything should be the supreme law
|name             =University of Mannheim
|native_name      =Universität Mannheim
|image_name       =Uni_Mannheim_Siegel.gif
|caption          =[[Seal (emblem)|Seal]] of the UMA
|established      =1763: Theodoro Palatinae <br/> 1907: Handelshochschu
|type             =[[Public University|Public]]
|endowment        =€115 [[million]]
|academic_staff   =800 (full time)
|administrative_staff  = 550 (full time)
|Schools          =5
|rector           =[[Ernst-Ludwig von Thadden]]
|chancellor       =[[Susann-Annette Storm]]
|students         =12,151 <small>''(HWS 2013/14)''</small><ref name="ur
/Studierendenstatistik_hws13.pdf|title= Studierendenstatistik der Unit
|undergrad        =6,915<ref name="uni-mannheim.de"/>
|postgrad         =4,965<ref name="uni-mannheim.de"/>
|doctoral         =249<ref name="uni-mannheim.de"/>
|profess          =
|city             =[[Mannheim]]
|state            =[[Baden-Württemberg]]
|country          =[[Germany]]
|coor             = {{Coord|49.4832|8.4647|region:DE-BW_type:edu_source
```

```
– <rdf:RDF>
  – <rdf:Description rdf:about="http://dbpedia.org/resource/Mannheim_Centre_for_European_Social_Research">
      //dbpedia.org/resource/University_of_Mannheim"/>

      edia.org/resource/Wolfgang_Franz">
      ://dbpedia.org/resource/University_of_Mannheim"/>
      /dbpedia.org/resource/University_of_Mannheim"/>
      /dbpedia.org/resource/University_of_Mannheim"/>

      edia.org/resource/Heinz_K%C3%B6nig">
      ://dbpedia.org/resource/University_of_Mannheim"/>

      edia.org/resource/Roman_Inderst">
      ://dbpedia.org/resource/University_of_Mannheim"/>
      ://dbpedia.org/resource/University_of_Mannheim"/>

      edia.org/resource/Claus_E._Heinrich">
      //dbpedia.org/resource/University_of_Mannheim"/>
      //dbpedia.org/resource/University_of_Mannheim"/>
      ://dbpedia.org/resource/University_of_Mannheim"/>

      edia.org/resource/Susann-Annette_Storm">
      ://dbpedia.org/resource/University_of_Mannheim"/>

      edia.org/resource/Bruno_Sälzer">
      ://dbpedia.org/resource/University_of_Mannheim"/>

  – <rdf:Description rdf:about="http://dbpedia.org/resource/Heinz_König">
      <dbo:award rdf:resource="http://dbpedia.org/resource/University_of_Mannheim"/>
    </rdf:Description>
```

DBpedia

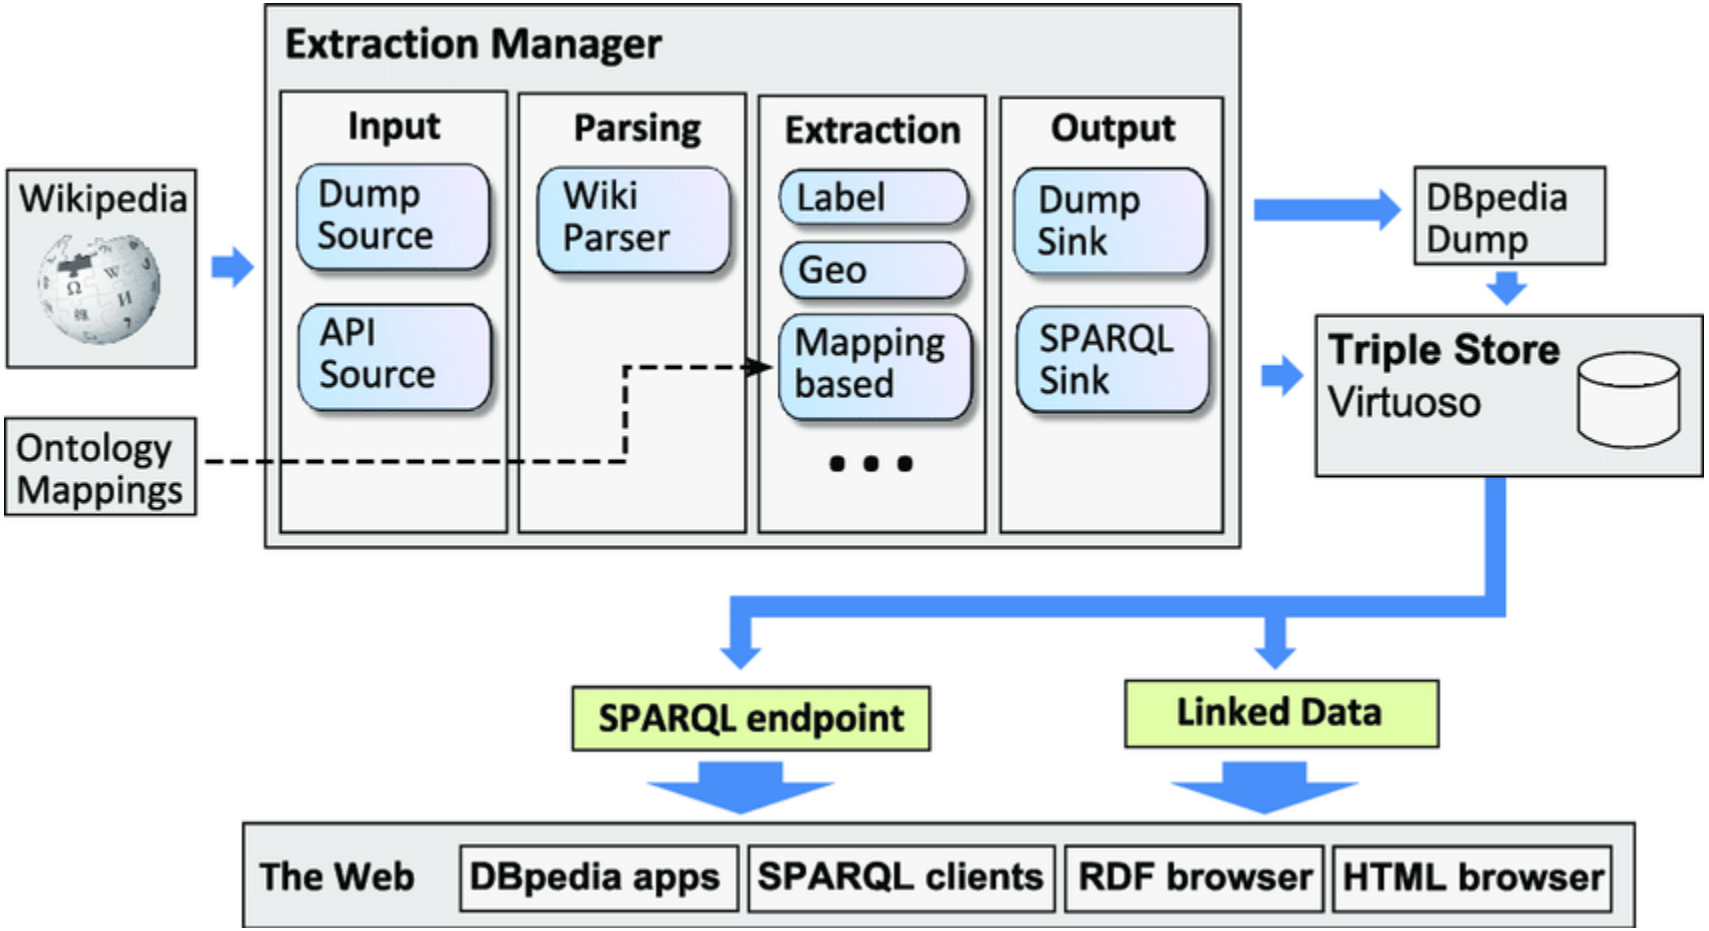

Lehmann et al.: *DBpedia – A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia.* 2014

DBpedia

Mapping en:Infobox film

This is the mapping for the Wikipedia template Infobox film. Find usages of this Wikipedia template here.

Test this mapping (or in namespace File or Creator) with some example Wikipedia pages. Check which prop

Read more about mapping Wikipedia templates.

Template Mapping (help)	
map to class	Film

Mappings

Property Mapping (help)	
template property	director
ontology property	director

Property Mapping (help)	
template property	producer
ontology property	producer

OntologyClass:Film

This is the definition of an ontology class.

Show all properties available for this class.

Show class in class hierarchy.

Read more about editing the ontology schema.

You can see the result of your edit on DBpedia Live (this is

Ontology class (help)	
rdfs:label (en)	film
rdfs:label (en)	movie
rdfs:label (nl)	film
rdfs:label (da)	film
rdfs:label (de)	Film
rdfs:label (el)	ταινία
rdfs:label (fr)	film
rdfs:label (ko)	영화
rdfs:label (ja)	映画
rdfs:label (ar)	فيلم
rdfs:label (pl)	film
rdfs:label (ga)	scannán
rdfs:label (es)	película
rdfs:subClassOf	Work
owl:equivalentClass	schema:Movie, wikidata:Q11424
owl:disjointWith	

OntologyProperty:director

This is the definition of an ontology property.

Read more about editing the ontology schema.

You can see the result of your edit on DBpedia Live (this is BETA!).

Ontology object property (help)	
rdfs:label (en)	director
rdfs:label (en)	film director
rdfs:label (nl)	regisseur
rdfs:label (da)	instruktør
rdfs:label (de)	regisseur
rdfs:label (ru)	директор
rdfs:label (el)	σκηνοθέτης
rdfs:label (es)	director de cine
rdfs:label (fr)	réalisateur
rdfs:comment (en)	A film director is a person who directs the making of a film.[1]
rdfs:comment (fr)	Un réalisateur (au féminin, réalisatrice) est une personne qui dirige la fabrication d'une œuvre audi cinéma ou la télévision.[2]
rdfs:domain	Film
rdfs:range	Person
rdf:type	
rdfs:subPropertyOf	dul:coparticipatesWith
owl:equivalentProperty	schema:director, wikidata:P57
owl:propertyDisjointWith	

DBpedia

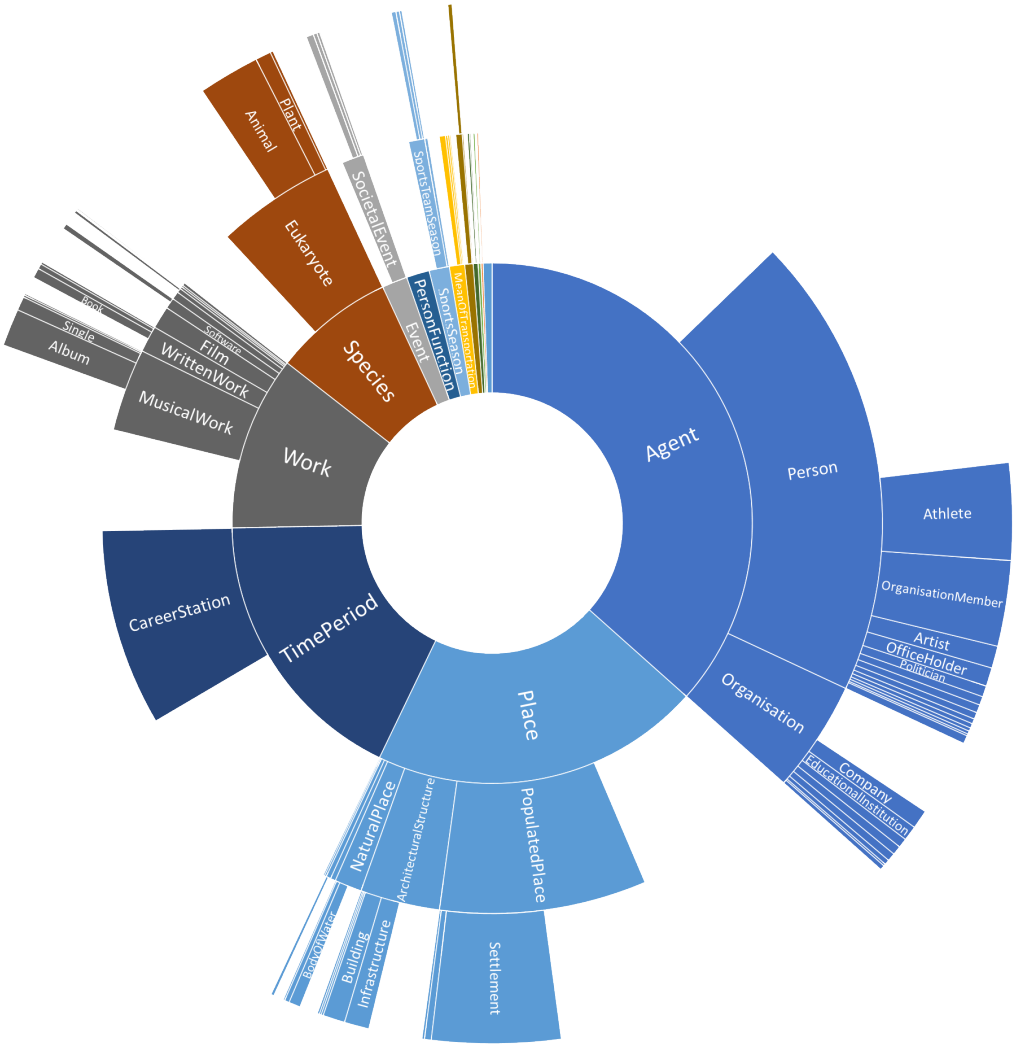
Agent
Person
Athlete
OrganisationMember
Artist
OfficeHolder
Politician
Organisation
Company
EducationalInstitution
Place
PopulatedPlace
Settlement
ArchitecturalStructure
Building
Infrastructure
BodyOfWater
NaturalPlace
TimePeriod
CareerStation
Work
MusicalWork
WrittenWork
Film
Software
Album
Single
Book
Species
Eukaryote
Animal
Plant
SocietalEvent
Event
PersonFunction
SportsSeason
SportsTeamSeason
MeanOfTransportation

YAGO

- Wikipedia categories for types
  - Plus WordNet as upper structure
- Manual mappings for properties

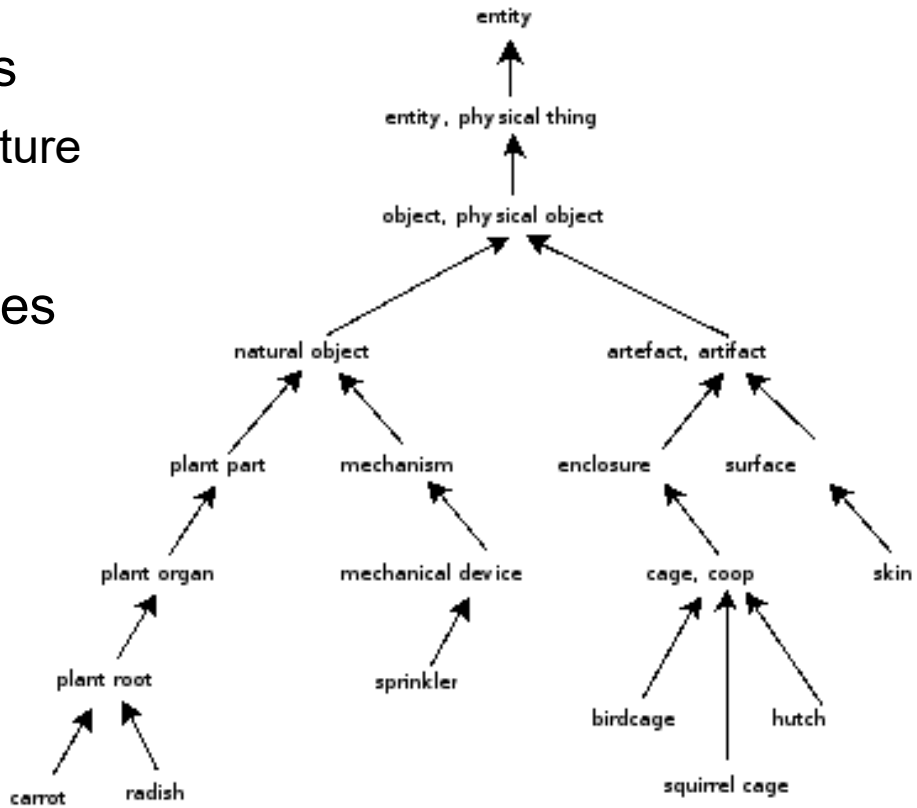

Figure 1. "is a" relation example

https://www.cs.princeton.edu/courses/archive/spring07/cos226/assignments/wordnet.html

YAGO

Datei Bearbeiten Ansicht Chronik Lesezeichen Extras Hilfe

Max-Planck-Institut für Infor... | YAGO Text Browser

https://gate.d5.mpi-inf.mpg.de/webyago3spotlx/Browser?entity=<Mannheim>

Suchen

<Mannheim>

← <Reinhard_Bütikofer>
← <Ümit_Davala>
← <Hans_Martin_Pippart>
← <de/Klaus_May>
← <Werner_Catel>
← <de/Fritz_Rößling>
← <Reinhold_Fanz>
← <Peter_Dreher>
← <Hans-Jürgen_Boysen>
← <Albert_Speer>
← <Manuel_Gulde>
← <Caroline_Augusta_of_Bavaria>
← <Reiner_Hollich>

<hasCitationTitle>

"Metropolregion Rhein-Neckar"@en
"World's 15 Most Inventive Cities"@
"The Manhattan of Germany: the in
"The rise of the smart city"@eng
"Germany and the Second World W
"Ausgabe der Klimadaten: Monatsv
"Press release announcing the mer
"FEI European Jumping Champions
"Partner und Freundesstädte"@eng
"Oraşe înfrăţite (Twin cities of Minsl
"Swansea - Wales :Mannheim.de"@

<Sepp_Herberger>

Search: eng

<id_1u5xrvs_1ul_zxcbb2>

<Miroslav_Klose> <playsFor> <FC_Bayern_Munich> hasFactId

<extractionSource> <http://en.wikipedia.org/wiki/Miroslav_Klose> →
<http://en.wikipedia.org/wiki/Miroslav_Klose> →

<occursUntil> "2011-##-##"^^xsd:date →

<occursSince> "2007-##-##"^^xsd:date →

<Lithuania>

YAGO


physical entity
object
whole
wikicat_Objects
wikicat_Physical_objects
location
wikicat_Artificial_objects
causal agent
person
abstraction
group
social group
communication
event
psychological feature
legal actor or geographic entity
legal actor
organization
person
location
region
yago Permanently Located Entity
organization
wikicat_Organizations
geographical entity
event
social event
structure
location

Knowledge Graph Creation

- Lesson learned no. 4:
  - Heuristics help increasing coverage (at the cost of accuracy)

Knowledge Graph Creation: NELL

- The 2010s
  - NELL: Never ending language learner
  - Input: ontology, seed examples, text corpus
  - Output: facts, text patterns
  - Large degree of automation, occasional human feedback

- Until 2018
  - Continuously ran for ~8 years
  - New release every few days

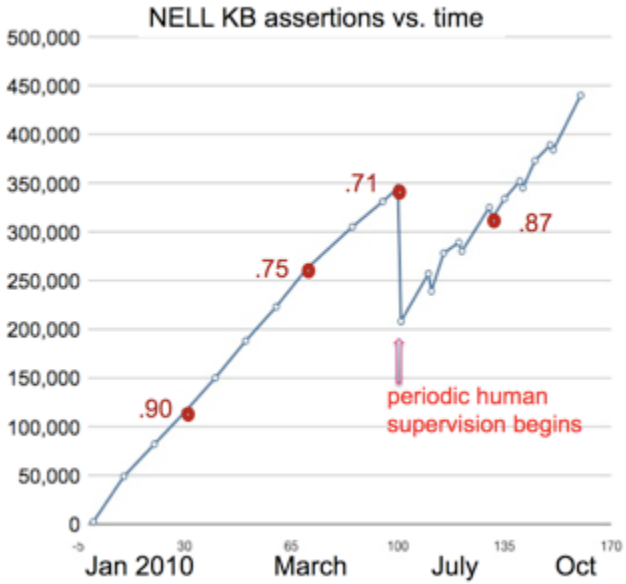

http://rtw.ml.cmu.edu/rtw/overview

Knowledge Graph Creation: NELL

- Extraction of a Knowledge Graph from a Text Corpus

Nine Inch Nails singer Trent Reznor, bor 196

...as stated by Filter singer Richard Pa

...says Slipknot singer Corey Taylo 44, in the interview.

patterns

"X singer Y"
→ band_member(X,Y)

facts

band_member(Nine_Inch_Nails, Trent_Reznor)
d_member(Filter,Richard_Patrick)
band_member(Slipknot,Corey_Taylor)

Recently-Learned Facts twitter Refresh

instance	iteration	date learned	confidence
conversion_table is an item found on the floor	1111	06-jul-2018	99.8
arlene_martel is a commedian	1111	06-jul-2018	94.0
cigar_rights is a socio-political term	1111	06-jul-2018	95.7
linton_zoological_gardens is an aquarium	1111	06-jul-2018	100.0
robb_miller coaches a sports team	1111	06-jul-2018	91.4
eric_e__schmidt is a person who was written about in new_york_times	1111	06-jul-2018	100.0
rodin is a visual artist in the field of sculpture	1115	03-sep-2018	99.6
the_today_show is a company in the economic sector of news	1114	25-aug-2018	93.0
china is a country located in the geopolitical location other_countries	1111	06-jul-2018	100.0
jerusalem is a city located in the geopolitical location israel	1114	25-aug-2018	99.8

Knowledge Graph Creation: NELL

Knowledge Graph Creation

- Lesson learned no. 5:
  - Quality cannot be maximized without human intervention

Summary of Trade Offs

- (Manual) effort vs. accuracy and completeness
- User involvement (or usability) vs. degree of formality
- Simplicity vs. support for plurality and provenance

→ all those decisions influence the shape of a knowledge graph!

Non-Public Knowledge Graphs

- Many companies have their own private knowledge graphs
  - Google: Knowledge Graph, Knowledge Vault
  - Yahoo!: Knowledge Graph
  - Microsoft: Satori
  - Facebook: Entities Graph
  - Thomson Reuters: permid.org (partly public)

- However, we usually know only little about them

See: Noy et al. (2019): Industry-scale Knowledge Graphs: Lessons and Challenges: Five diverse technology companies show how it's done

Comparison of Knowledge Graphs

- Release cycles

Instant updates:
DBpedia live,
Freebase
Wikidata

Days:
NELL

Caution!

Months:
DBpedia

Years:
YAGO
Cyc

- Size and density

Table 1: Global Properties of the Knowledge Graphs compared in this paper

	DBpedia	YAGO	Wikidata	OpenCyc	NELL
Version	2016-04	YAGO3	2016-08-01	2016-09-05	08m.995
# instances	5,109,890	5,130,031	17,581,152	118,125	1,974,297
# axioms	397,831,457	1,435,808,056	1,633,309,138	2,413,894	3,402,971
avg. indegree	13.52	17.44	9.83	10.03	5.33
avg. outdegree	47.55	101.86	41.25	9.23	1.25
# classes	754	576,331	30,765	116,822	290
# relations	3,555	93,659	11,053	165	1,334

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Comparison of Knowledge Graphs

- What do they actually contain?
- Experiment: pick 25 classes of interest
  - And find them in respective ontologies
- Count instances (coverage)
- Determine in and out degree (level of detail)

Comparison of Knowledge Graphs

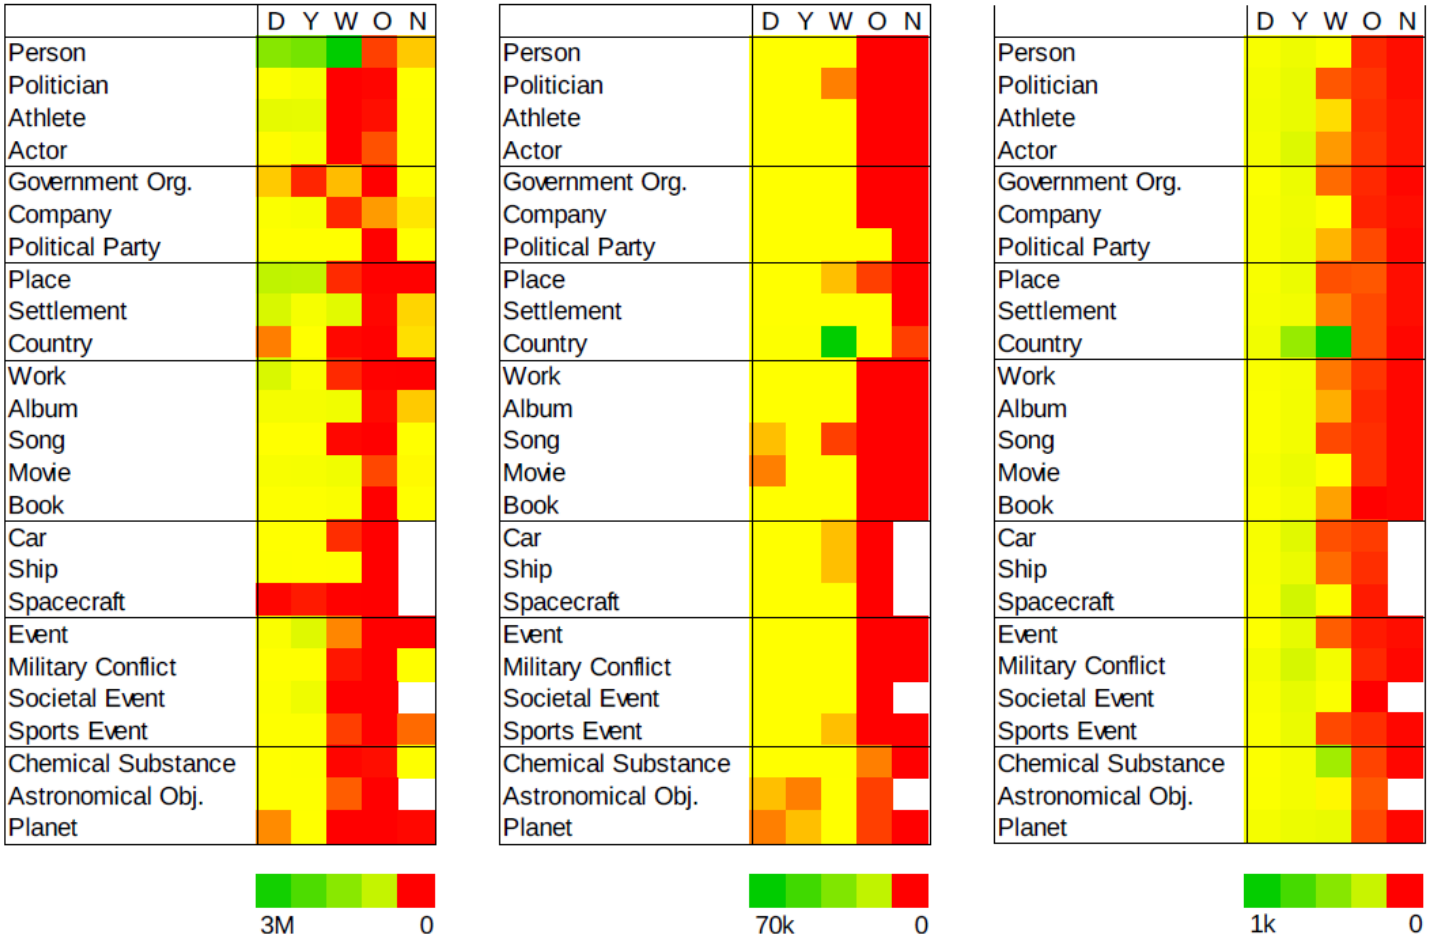

(a) Number of instances (b) Average indegree (c) Average outdegree

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Comparison of Knowledge Graphs

- Summary findings:
  - Persons: more in Wikidata
    (twice as many persons as DBpedia and YAGO)
  - Countries: more details in Wikidata
  - Places: most in DBpedia
  - Organizations: most in YAGO
  - Events: most in YAGO
  - Artistic works:
    - Wikidata contains more movies and albums
    - YAGO contains more songs

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Caveats

- Reading the diagrams right…

- So, Wikidata contains more persons
  - but less instances of all the interesting subclasses?

- There are classes like *Actor* in Wikidata
  - but they are hardly used
  - rather: modeled using *profession* relation

Caveats

- Reading the diagrams right… (ctd.)

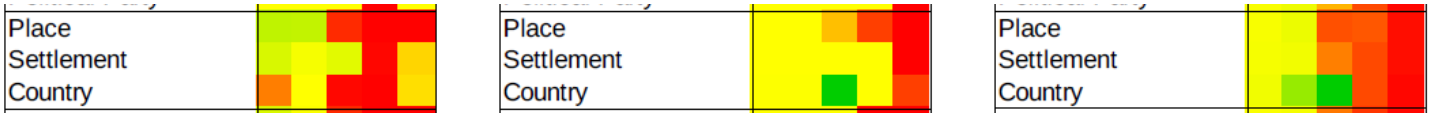

- So, Wikidata contains more data on countries, but less countries?
- First: Wikidata only counts current, actual countries
  - DBpedia and YAGO also count historical countries
- “KG1 contains less of X than KG2” can mean
  - it actually contains less instances of X
  - it contains equally many or more instances,
    but they are not typed with X (see later)
- Second: we count single facts about countries
  - Wikidata records some time indexed information, e.g., population
  - Each point in time contributes a fact

Overlap of Knowledge Graphs

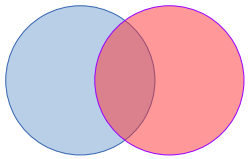

- How largely do knowledge graphs overlap?
- They are interlinked, so we can simply count links
  - For NELL, we use links to Wikipedia as a proxy

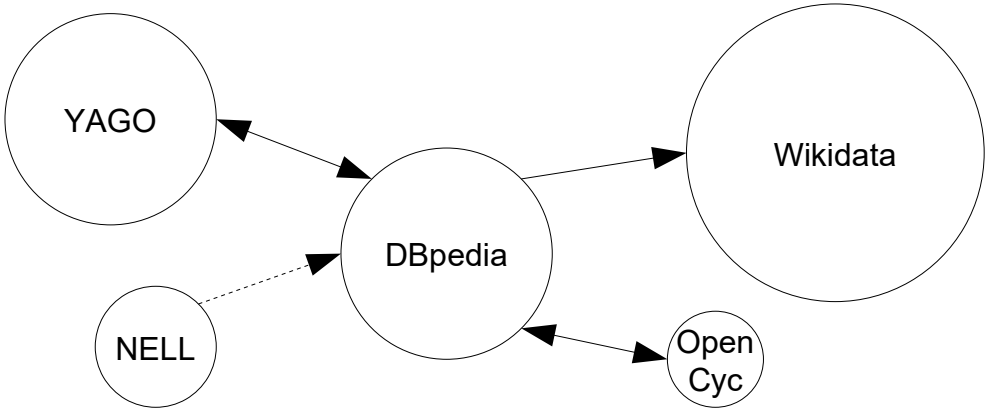

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Overlap of Knowledge Graphs

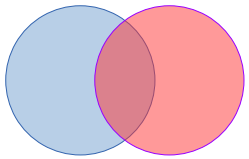

- How largely do knowledge graphs overlap?
- They are interlinked, so we can simply count links
  - For NELL, we use links to Wikipedia as a proxy

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Overlap of Knowledge Graphs

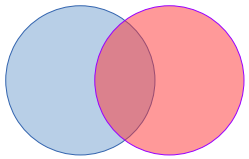

- Links between Knowledge Graphs are incomplete
  - The Open World Assumption also holds for interlinks

- But we can estimate their number
- Approach:
  - find link set automatically with different heuristics
  - determine precision and recall on existing interlinks
  - estimate actual number of links

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Overlap of Knowledge Graphs

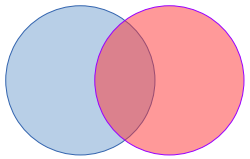

- Idea:
  - Given that the link set F is found
  - And the (unknown) actual link set would be C

- Precision P: Fraction of F which is actually correct
  - i.e., measures how much |F| is *over*-estimating |C|
- Recall R: Fraction of C which is contained in F
  - i.e., measures how much |F| is *under*-estimating |C|

- From that, we estimate $|C| = |F| \cdot P \cdot \frac{1}{R}$

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Overlap of Knowledge Graphs

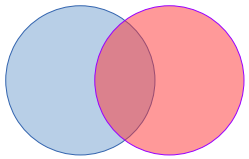

- Mathematical derivation:
  - Definition of recall: $R = \frac{|F_{correct}|}{|C|}$

    unknown
  - Definition of precision: $P = \frac{|F_{correct}|}{|F|}$
- Resolve both to $|F_{correct}|$, substitute, and resolve to $|C|$

$$|C| = |F| \cdot P \cdot \frac{1}{R}$$

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Overlap of Knowledge Graphs

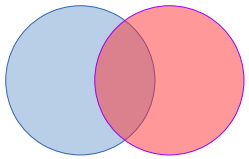

- Experiment:
  - We use the same 25 classes as before
  - Measure 1: overlap relative to smaller KG (i.e., potential gain)
  - Measure 2: overlap relative to explicit links (i.e., importance of improving links)

- Link generation with 16 different metrics and thresholds
  - Intra-class correlation coefficient for |C|: 0.969
  - Intra-class correlation coefficient for |F|: 0.646
- Bottom line:
  - Despite variety in link sets generated, the overlap is estimated reliably
  - The link generation mechanisms do not need to be overly accurate

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Overlap of Knowledge Graphs

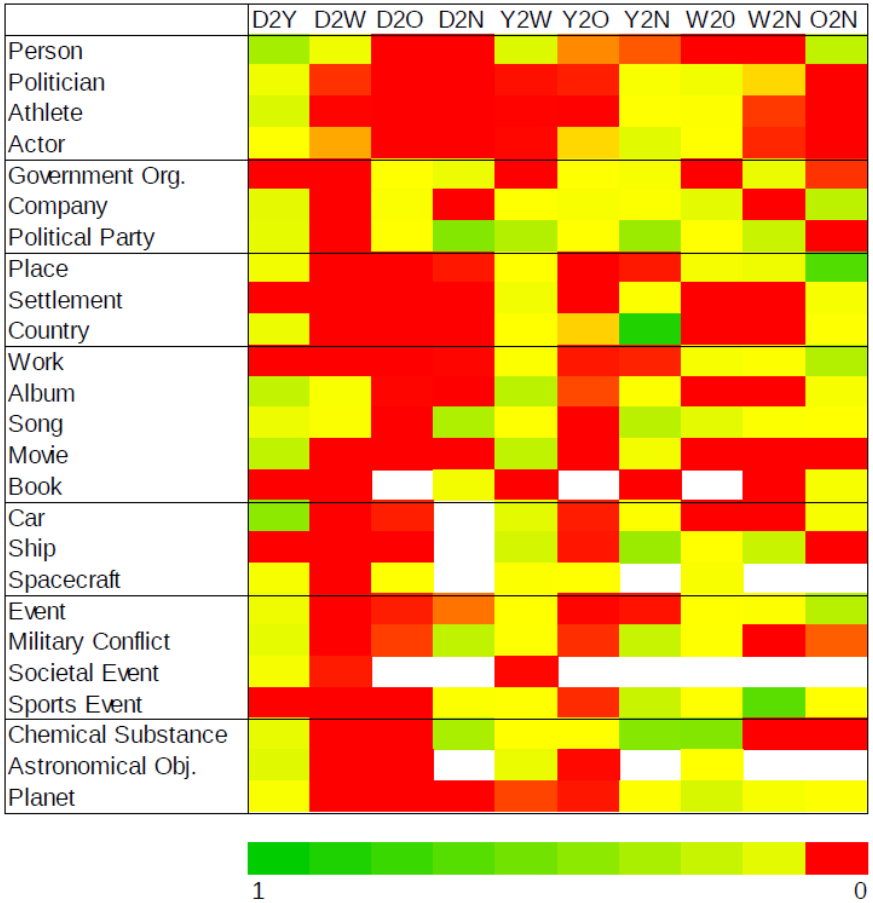

(a) Overlap as potential gain

(b) Overlap relative to existing links

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Overlap of Knowledge Graphs

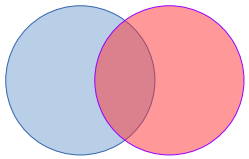

- Summary findings:
  - DBpedia and YAGO cover roughly the same instances (not much surprising)
  - NELL is the most complementary to the others
  - Existing interlinks are insufficient for out-of-the-box parallel usage

Ringler & Paulheim: *One Knowledge Graph to Rule them All?* KI 2017

Intermezzo: Knowledge Graph Creation Cost

- There are quite a few metrics for evaluating KGs
  - size, degree, interlinking, quality, licensing, ...

Table 2

Data quality metrics related to accessibility dimensions (type QN refers to a quantitative metric, QL to a qualitative one).

Dimension	Abr	Metric	Description	Type
Availability	A1	accessibility of the SPARQL end-point and the server	checking whether the server responds to a SPARQL query [18]	QN
	A2	accessibility of the RDF dumps	checking whether an RDF dump is provided and can be downloaded [18]	QN
	A3	dereferenceability of the URI	checking (i) for dead or broken links i.e. when an HTTP-GET request is sent, the status code `404 Not Found` is not be returned (ii) that useful data (particularly RDF) is returned upon lookup of a URI, (iii) for changes in the URI i.e the compliance with the recommended way of implementing redirections using the status code `303 See Other` [18,30]	QN
	A4	no misreported content types	detect whether the HTTP response contains the header field stating the appropriate content type of the returned file e.g. `application/rdf+xml` [30]	QN
	A5	dereferenced forward-links	dereferenceability of all forward links: all available triples where the local URI is mentioned in the subject (i.e. the description of the resource) [31]	QN
Licensing	L1	machine-readable indication of a license	detection of the indication of a license in the VoID description or in the dataset itself [18,31]	QN
	L2	human-readable indication of a license	detection of a license in the documentation of the dataset [18, 31]	QN
	L3	specifying the correct license	detection of whether the dataset is attributed under the same license as the original [18]	QN
Interlinking	I1	detection of good quality interlinks	(i) detection of (a) interlinking degree, (b) clustering coefficient, (c) centrality, (d) open sameAs chains and (c) description richness through sameAs by using network measures [25], (ii) via crowdsourcing [1,65]	QN
	I2	existence of links to external data providers	detection of the existence and usage of external URIs (e.g. using `owl:sameAs` links) [31]	QN
	I3	dereferenced back-links	detection of all local in-links or back-links: all triples from a dataset that have the resource's URI as the object [31]	QN
Security	S1	usage of digital signatures	by signing a document containing an RDF serialization, a SPARQL result set or signing an RDF graph [13,18]	QN
	S2	authenticity of the dataset	verifying authenticity of the dataset based on a provenance vocabulary such as author and his contributors, the publisher of the data and its sources (if present in the dataset) [18]	QL
Performance	P1	usage of slash-URIs	checking for usage of slash-URIs where large amounts of data is provided [18]	QN
	P2	low latency	(minimum) delay between submission of a request by the user and reception of the response from the system [18]	QN
	P3	high throughput	(maximum) no. of answered HTTP-requests per second [18]	QN
	P4	scalability of a data source	detection of whether the time to answer an amount of ten requests divided by ten is not longer than the time it takes to answer one request [18]	QN

Zaveri et al.: *Quality Assessment for Linked Open Data: A Survey.* SWJ 7(1), 2016

Table 14

Framework with an example weighting which would be reasonable for a user setting as given in [30].

Dimension	Metric	DBpedia	Freebase	OpenCyc	Wikidata	YAGO	Example of User Weighting w_i
Accuracy	m_{synRDF}	1	1	1	1	1	1
	m_{synLit}	0.994	1	1	1	0.624	1
	$m_{semTriple}$	1	1	1	1	1	1
Trustworthiness	m_{graph}	0.5	0.5	1	0.75	0.25	1
	m_{fact}	0.5	1	0	1	1	2
	m_{NoVal}	0	1	0	1	0	1
Consistency	$m_{checkRestr}$	0	1	0	1	0	1
	$m_{conClass}$	0.875	1	0.999	1	0.333	1
	$m_{conRelat}$	0.991	0.45	1	0	0.992	1
Relevancy	$m_{Ranking}$	0	0	0	1	0	1
Completeness	$m_{cSchema}$	0.905	0.762	0.921	1	0.952	1
	m_{cCol}	0.402	0.425	0	0.285	0.332	1
	m_{cPop}	0.93	0.94	0.48	0.99	0.89	3
Timeliness	m_{Freq}	0.5	0	0.25	1	0.25	3
	$m_{Validity}$	0	1	0	1	1	1
	m_{Change}	0	1	0	0	0	1
Ease of understanding	m_{Descr}	0.704	0.972	1	0.9999	1	3
	m_{Lang}	1	1	0	1	1	2
	m_{uSer}	1	1	0	1	1	1
	m_{uURI}	1	0.5	1	0	1	2
Interoperability	m_{Reif}	1	0.5	0.5	0	0.5	1
	$m_{iSerial}$	1	0	0.5	1	1	2
	m_{extVoc}	0.61	0.108	0.415	0.682	0.134	2
	$m_{propVoc}$	0.15	0	0.513	0.001	0	1
Accessibility	m_{Deref}	1	0.437	1	0.414	1	2
	m_{Avai}	0.9961	0.9998	1	0.9999	0.7306	2
	m_{SPARQL}	1	0	0	1	1	1
	m_{Export}	1	1	1	1	1	0
	m_{Negot}	0.5	0	0	1	1	1
	m_{HTML_RDF}	1	1	0	1	1	0
	m_{Meta}	1	0	1	0	0	1
Licensing	$m_{machLicense}$	1	0	0	1	0	1
Interlinking	m_{Inst}	0.592	0.018	0.443	0	0.305	2
	m_{URIs}	0.929	0.954	0.894	0.957	0.956	1
Unweighted Average		0.708	0.605	0.498	0.738	0.625	
Weighted Average		0.718	0.575	0.516	0.742	0.646	

Färber et al.: *Linked data quality of DBpedia, Freebase, OpenCyc, Wikidata, and YAGO* SWJ 9(1), 2018

Intermezzo: Knowledge Graph Creation Cost

- ...but what is the cost of a single statement?

Some back of the envelope calculations...

Paulheim: How much is a triple?
Estimating the Cost of Knowledge Graph Creation, 2018

Intermezzo: Knowledge Graph Creation Cost

- Case 1: manual curation
  - Cyc: created by experts
    Total development cost: $120M
    Total #statements: 21M

    → **$5.71 per statement**
  - Freebase: created by laymen
    Assumption: adding a statement to Freebase equals adding a sentence to Wikipedia
    - English Wikipedia up to April 2011: 41M working hours (Geiger and Halfaker, 2013),

      size in April 2011: 3.6M pages, avg. 36.4 sentences each
    - Using US minimum wage: $2.25 per sentence

    → **$2.25 per statement**

    (Footnote: total cost of creating Freebase would be $6.75B)

acquisition by Google estimated as $60-300M

Intermezzo: Knowledge Graph Creation Cost

- Case 2: automatic/heuristic creation
  - DBpedia: 4.9M LOC, 2.2M LOC for mappings

    software project development: ~37 LOC per hour (Devanbu et al., 1996)

    we use German PhD salaries as a cost estimate

    → **1.85c per statement**
  - YAGO: made from 1.6M LOC

    uses WordNet: 117k synsets, we treat each synset like a Wiki page

    → **0.83c per statement**
  - NELL: 103k LOC

    → **14.25c per statement**
- Compared to manual curation: saving factor 16-250

Intermezzo: Knowledge Graph Creation Cost

- Graph error rate against cost
  - we can pay for accuracy
  - NELL is a bit of an outlier

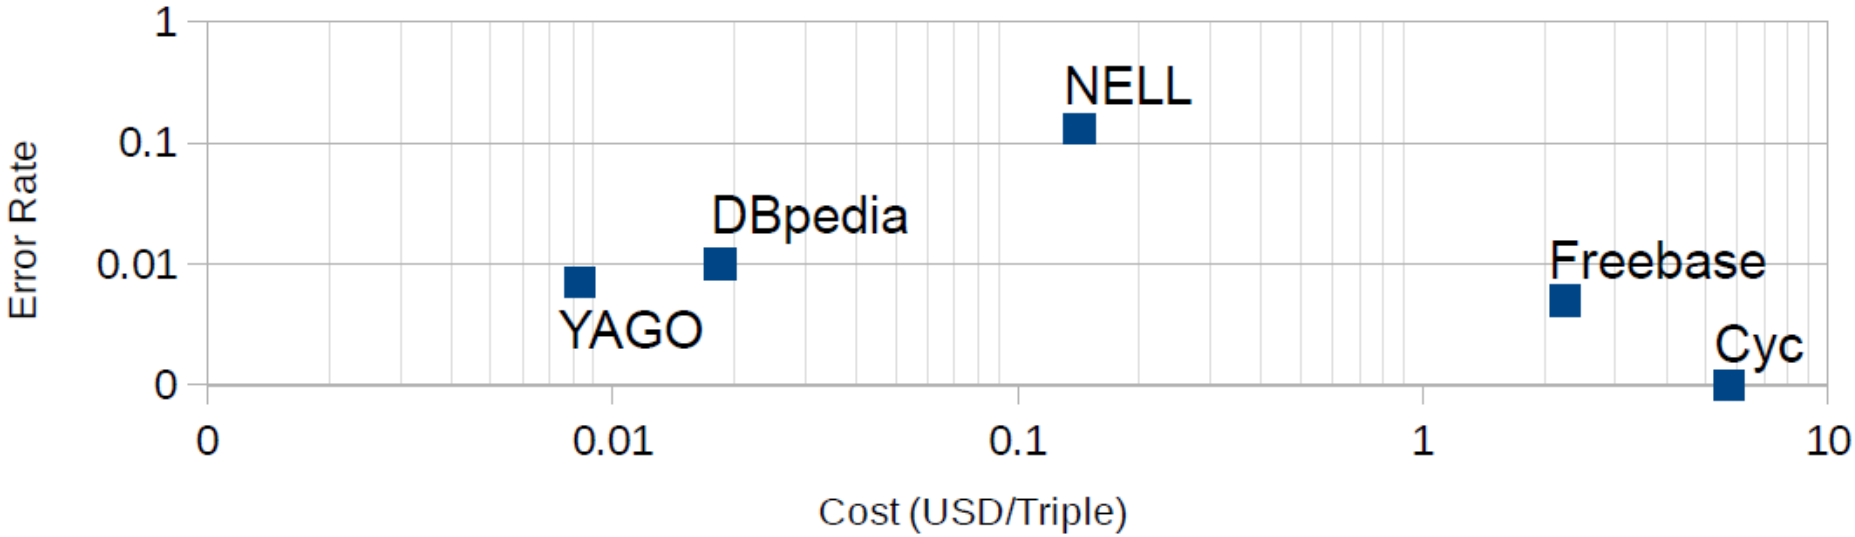

Enhancing the Coverage of Knowledge Graphs

- Study for KG-based Recommender Systems*
  - DBpedia (likewise: YAGO) has a coverage of
    - 85% for movies
    - 63% for music artists
    - 31% for books

Delicious Bookmarks

105,000 bookmarks from 1867 users.

- README.txt
- hetrec2011-delicious-2k.zip

Last.FM

92,800 artist listening records from 1892 users.

- README.txt
- hetrec2011-lastfm-2k.zip

MovieLens + IMDb/Rotten Tomatoes

86,000 ratings from 2113 users.

- README.txt
- hetrec2011-movielens-2k.zip

https://grouplens.org/datasets/

*) Di Noia, et al.: *SPRank: Semantic Path-based Ranking for Top-n Recommendations using Linked Open Data.* In: ACM TIST, 2016

Enhancing the Coverage of Knowledge Graphs

- Only existing pages have categories
  - Lists may also link to non-existing pages

List of intelligent dance music artists

From Wikipedia, the free encyclopedia

This section **does not cite any sources**. Please help improve this section by adding citations to reliable sources. Unsourced
Find sources: "List of intelligent dance music artists" – news · newspapers · books · scholar · JSTOR (June 2015) (Learn how and when to remov

This is a list of notable music artists who play intelligent dance music (IDM) genre.

Contents [hide]

[edit]

- 808 State
- µ-ziq

A-K [edit]

- Actress
- Acustic
- Air Liquide
- Alarm Will Sound
- Alva Noto
- Amon Tobin
- Andy Stott
- Aphex Twin
- Apparat
- Arovane
- Atypic
- Autechre
- B12
- Benn Jordan
- Biosphere
- Björk[1]
- The Black Dog
- Blanck Mass
- Boards of Canada
- Bochum Welt
- Boom Bip
- Brothomstates
- Burial
- Bvdub
- C418
- Cabaret Voltaire
- Casino Versus Japan
- Ceephax Acid Crew
- Cex
- Christ.
- Chris Clark
- Clocolan
- Cylob
- Daedelus
- Deadbeat
- Deepchord
- Demdike Stare
- Deru
- Richard Devine
- Dopplereffekt
- Chris Douglas
- Drexciya
- Eight Frozen Modules
- Emptyset
- Esem
- FaltyDL
- Fennesz
- The Field
- The Flashbulb
- Floating Points
- Flying Lotus
- Forest Swords
- Funkstörung
- The Future Sound Of London
- Gas
- Gescom
- Global Communication
- Global Goon
- Goldie
- Zachary Gray[2]
- Gridlock
- Himuro Yoshiteru
- Kim Hiorthøy
- I am Robot and Proud
- Innovaders
- Jan Jelinek
- Jega
- Jello
- Jlin
- John Tejada
- Jon Hopkins[3]
- Kettel
- Kevin Blechdom
- Kid606
- Kodomo
- Koreless[4]

Delicious Bookmarks

105,000 bookmarks from 1867 users.

- README.txt
- hetrec2011-delicious-2k.zip

Last.FM

92,800 artist listening records from 1892 users.

- README.txt
- hetrec2011-lastfm-2k.zip

MovieLens + IMDb/Rotten Tomatoes

86,000 ratings from 2113 users.

- README.txt
- hetrec2011-movielens-2k.zip

Entity Extraction from List Pages

- Lists form (shallow) hierarchies

Lists of writers

From Wikipedia, the free encyclopedia

The following are **lists of writers**:

- Bestsellers
- Biographers
- Buddhism
- Business theorists
- Catholicism
- Children's literature
- Christian fiction
- Cricket
- Crime
- Detective fiction
- Drama
- Essays
- Fantasy
- History
- Historical novels
- Horror fiction
- Horsemanship
- Illustrations
- Manga
- Music theory
- Mysteries
- Non-fiction
- Novels
- Occult
- Plays
- Poetry
- Politics
- Role-playing games
- Romantic novels
- Science fiction
- Self-help
- Short stories
- Software
- Technical writers
- Thrillers
- Translations
- Western fiction
- Young adult

Top of page

Contents [hide]

Lists by language (non-English) [edit]

- Ancient Greek
- Arabic
- Bengali
- Catalan
- Chichewa
- Dutch
- French
- German
- Gujarati
- Greek
- Hebrew
- Hindi
- Leonese
- Lithuanian
- Malayalam
- Marathi
- Nepali
- Odia
- Pukhto or Pashto
- Persian
- Polish
- Portuguese
- Russian
- Spanish
- Swedish
- Tamil
- Turkish
- Urdu
- Welsh

Top of page

Lists by ethnicity or nationality [edit]

- African writers
- African-American writers
- Albanian writers
- Algerian writers
- American writers
- Arab American writers
- Armenian authors
- Barbadian writers
- Beninese writers
- Black British writers
- Bosniak writers
- Brazilian writers
- Canadian writers
- Chinese writers
- English writers
- Egyptian writers
- Georgian writers
- German writers
- Ghanaian writers
- Greek writers
- Guyanese writers
- Irish writers
- Italian writers
- Jewish American writers
- Kenyan writers
- Korean American writers
- Macedonian writers
- Mexican writers
- Nepali writers
- New Zealand writers
- Nigerian writers
- Norwegian writers
- American novelists
- Pakistani writers
- Peruvian writers
- Romanian writers
- Russian authors
- Scottish writers
- Serbian writers
- Slovak authors
- Slovenian writers
- Somali writers

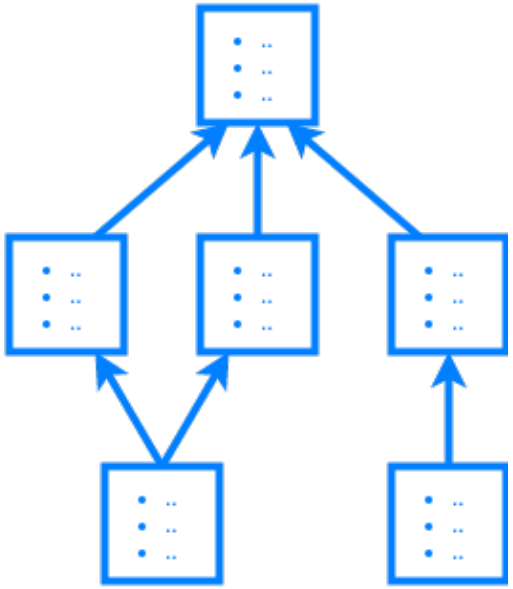

Entity Extraction from List Pages

- Idea: align with category graph
- Equivalence:
  - “List of Japanese Writers”
    ↔ Category:Japanese Writers
- Subsumption:
  - “List of Japanese
    Speculative Fiction Writers”
    → Category:Japanese Writers

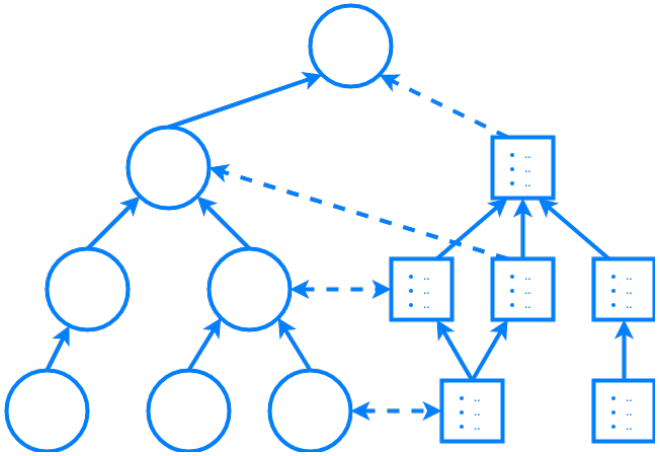

Classifying Red Links

- Not all entities on a list page belong to the same category
- Idea:
  - Learn classifier to tell subject entities from non-subject entities
- Distant learning approach
  - Positive examples:
    - Entities that are in the corresponding category
  - Negative examples
    - Entities that are in a category which is disjoint
    - e.g., Book <> Writer

- Patricia Aakhus (1952–2012), *The Voyage of Mael Duin's Curragh*
- Atia Abawi
- Edward Abbey (1927–1989), *The Monkey Wrench Gang*
- Lynn Abbey (born 1948), *Daughter of the Bright Moon*
- Belle Kendrick Abbott (1842–1893), *Leah Mordecai*
- Eleanor Hallowell Abbott (1872–1958), poet, novelist and short story writer
- Hailey Abbott, *Summer Boys*
- Megan Abbott (born 1971), *Die A Little*
- Shana Abé, *A Rose in Winter*
- Louise Abeita (1926–2014), Native American Isleta Pueblo writer, *I am a Pueblo Indian Girl*
- Robert H. Abel (1941–2017)
- Aberjhani
- Walter Abish (born 1931), *How German Is It*
- Abiola Abrams (born 1976), TV host, art filmmaker and author, *Dare*
- Diana Abu-Jaber (born 1960), *Arabian Jazz*
- Susan Abulhawa, *Mornings in Jenin*
- Kathy Acker (1947–1997), *Blood and Guts in High School*
- Cherry Adair, *Black Magic*
- Alice Adams (1926–1999), *Beautiful Girl*
- Victoria Aveyard (born 1990), *Red Queen* series

Increasing Level of Detail

- YAGO uses categories for types
  - e.g., Category:American Industrial Groups
  - but does not analyze them further

- `:NineInchNails a :AmericanIndustrialGroup`
  - “Things, not Strings”?

-
```
:NineInchNails  a :MusicalGroup ;
                hometown :United_States ;
                genre :Industrial .
```

Cat2Ax: Axiomatizing Wikipedia Categories

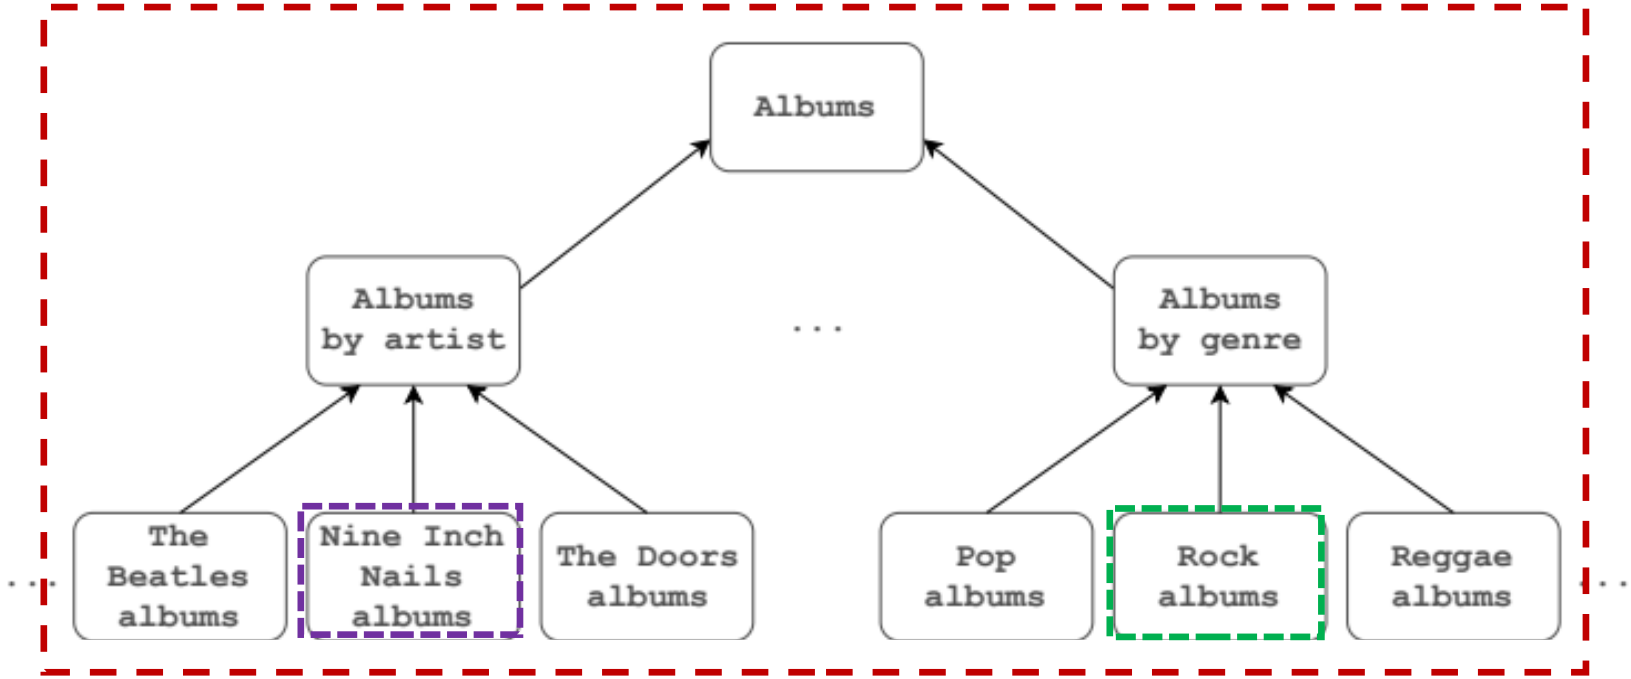

$\subseteq$ dbo:Album

$\subseteq$ dbo:artist.{dbr:Nine_Inch_Nails}

$\subseteq$ dbo:genre.{dbr:Rock_Music}

Heist & Paulheim (2019): Uncovering the Semantics of Wikipedia Categories

Cat2Ax: Axiomatizing Wikipedia Categories

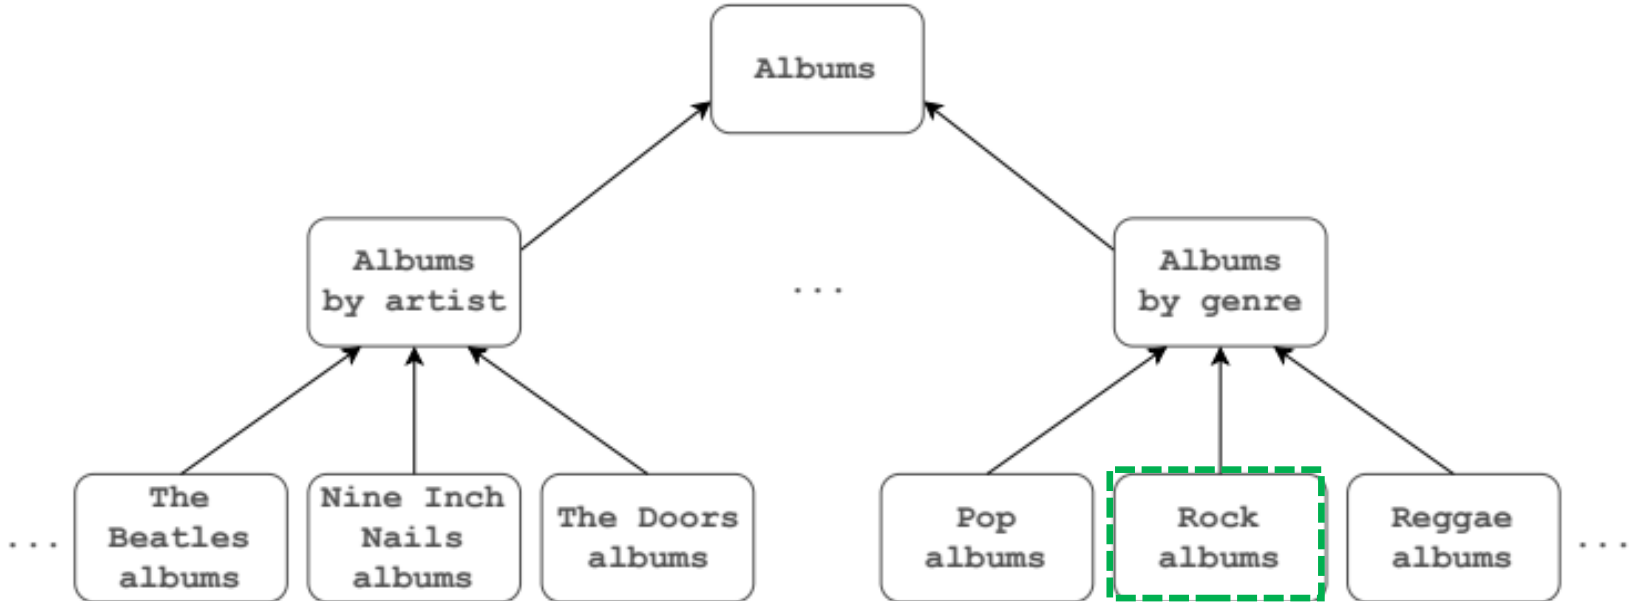

$\subseteq$ dbo:genre.{dbr:Rock_Music} ?

$\subseteq$ dbo:artist.{dbr:Rock_(Rapper)} ?

Cat2Ax: Axiomatizing Wikipedia Categories

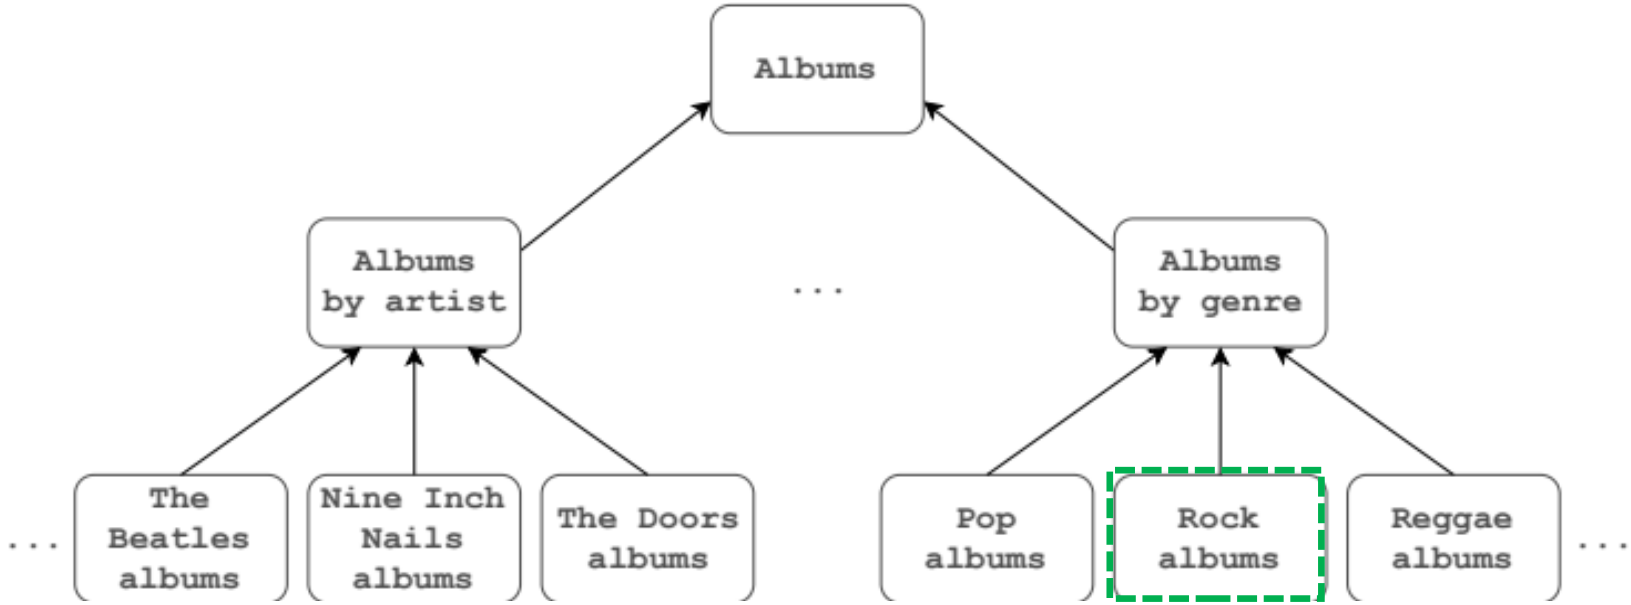

- Frequency: how often does the pattern occur in a category?
  - i.e.: share of instances that have dbo:genre.{dbr.Rock_Music}?
- Lexical score: likelihood of term as a surface form of object
  - i.e.: how often is *Rock* used to refer to dbr:Rock_Music?
- Sibling score: how likely are sibling categories sharing similar patterns?
  - i.e., are there sibling categories with a high score for dbo:genre?

CaLiGraph Example

rdfs:label	• Tiamat
owl:sameAs	• dbr:Tiamat_(band)
clgo:activeYearsStartYear	• 1987
clgo:genre	• Symphonic metal
clgo:hometown	• Sweden

Category: Musical Groups established in 1987

List of symphonic metal bands

Category: Swedish death metal bands
List of Swedes in Music

Pushing Entity Coverage Further

- Beyond red links (2020)

Cinematic films

Title	Running time	Year released	Notes
Amra Ekta Cinema Banabo (The Innocence)	1265 min (21 hr, 5 min)	2019	[31][32]
Resan (The Journey)	873 min (14 hr, 33 min)	1987	[33]
La Flor	803 min (13 hr, 23 min)	2018	[34]
Out 1 (Noli me tangere)	775 min (12 hr, 55 min)	1971	[35]
Evolution of a Filipino Family	593 min (9 hr, 53 min)	2004	[36]
Shoah	566 min (9 hr, 26 min)	1985	[37]
Tie Xi Qu: West of the Tracks	551 min (9 hr, 11 min)	2003	[38]
Death in the Land of Encantos	538 min (8 hr, 58 min)	2007	[39]
Dead Souls	495 min (8 hr, 15 min)	2018	[40]
A Lullaby to the Sorrowful Mystery	485 min (8 hr, 5 min)	2016	[41]
O.J.: Made in America	463 min (7 hr, 43 min)	2016	[42]
Melancholia	450 min (7 hr, 30 min)	2008	[43]
Sátántangó	419 min (6 hr, 59 min)	1994	[44]
La Roue	413 min (6 hr, 53 min)	1923 (Restoration, 2019)	[45]
The Best of Youth	366 min (6 hr, 6 min)	2003	[46]
Century of Birthing	360 min (6 hr)	2011	[47]
Near Death	358 min (5 hr, 58 min)	1989	[48]
Karamay	356 min (5 hr, 56 min)	2011	[49]
Little Dorrit	350 min (5 hr, 50 min)	1987	[50]
Carlos	339 min (5 hr, 39 min)	2010	[51]
Mula sa Kung Ano ang Noon	338 min (5 hr, 38 min)	2014	[52]
Napoléon	332 min (5 hr, 32 min)	1927 (Restoration, 2016)	[53]
1900	317 min (5 hr, 17 min)	1976	[54]
Happy Hour	317 min (5 hr, 17 min)	2015	[55]
Batang West Side	315 min (5 hr, 15 min)	2001	[56]
The Deluge	315 min (5 hr, 15 min)	1974	[57]
Fanny and Alexander	312 min (5 hr, 12 min)	1982	[58]
Tsahal	304 min (5 hr, 4 min)	1994	[59]

- Beyond explicit lists (2021)

Members [edit]

- Jürgen Engler – vocals, guitar, keyboards, synthesizers and programming, metallic percussion (1980–1985, 1989–1997, 2005–present)
- Ralf Dörper – keyboards, synthesizers and programming (1980–1982, 1985, 1989–1997, 2005–present)
- Marcel Zürcher – guitar, keyboards (2005–present)
- Nils Finkeisen - guitar (2015–present)
- Paul Keller - drums (2018–present)

Former members [edit]

- Bradley Bills - live drums (2013–2014)
- Rüdiger Esch - bass guitar (1989–1997, 2005–2011)
- Christoph "Nook" Michelfeit - drums, electronic percussion
- Bernward Malaka - bass guitar (1980–1982)
- Hendrik Thiesbrummel - live drums (2016–2018)
- Frank Köllges - drums
- Eva Gossling - saxophone (1981)
- Christina Schnekenburger - keyboards
- Walter Jäger - ?
- Christopher Lietz - programming, samples (1995–1997)
- Lee Altus - guitar (1992–1997)
- Darren Minter - drums (1993)
- George Lewis - drums (1997)
- Oliver Röhl – drums
- Achim Färber – drums
- Volker Borchert – drums (1992, 2015–2016)

Discography [edit]

Albums [edit]

- *Stahlwerksynfonie* (1981)
- *Volle Kraft Voraus!* (1982)
- *Entering the Arena* (1985)
- *I* (1992)
- *II - The Final Option* (1993)
- *The Final Remixes* (1994)
- *III - Odyssey of the Mind* (1995)
- *Paradise Now* (1997)
- *The Machinists of Joy* (2013)
- *V - Metal Machine Music* (2015)
- *Stahlwerkrequiem* (2016)
- *Live Im Schatten Der Ringe* (2016)

Entity Extraction from List Pages

- Red and grey links
  - Red links point to entities that do not exist
  - "Grey links"
    - are actually not links
    - i.e., entities to be discovered

Cinematic films

Title	Running time	Year released	Notes
Amra Ekta Cinema Banabo (The Innocence)	1265 min (21 hr, 5 min)	2019	[31][32]
Resan (The Journey)	873 min (14 hr, 33 min)	1987	[33]
La Flor	803 min (13 hr, 23 min)	2018	[34]
Out 1 (Noli me tangere)	775 min (12 hr, 55 min)	1971	[35]
Evolution of a Filipino Family	593 min (9 hr, 53 min)	2004	[36]
Shoah	566 min (9 hr, 26 min)	1985	[37]
Tie Xi Qu: West of the Tracks	551 min (9 hr, 11 min)	2003	[38]
Death in the Land of Encantos	538 min (8 hr, 58 min)	2007	[39]
Dead Souls	495 min (8 hr, 15 min)	2018	[40]
A Lullaby to the Sorrowful Mystery	485 min (8 hr, 5 min)	2016	[41]
O.J.: Made in America	463 min (7 hr, 43 min)	2016	[42]
Melancholia	450 min (7 hr, 30 min)	2008	[43]
Sátántangó	419 min (6 hr, 59 min)	1994	[44]
La Roue	413 min (6 hr, 53 min)	1923 (Restoration, 2019)	[45]
The Best of Youth	366 min (6 hr, 6 min)	2003	[46]
Century of Birthing	360 min (6 hr)	2011	[47]
Near Death	358 min (5 hr, 58 min)	1989	[48]
Karamay	356 min (5 hr, 56 min)	2011	[49]
Little Dorrit	350 min (5 hr, 50 min)	1987	[50]
Carlos	339 min (5 hr, 39 min)	2010	[51]
Mula sa Kung Ano ang Noon	338 min (5 hr, 38 min)	2014	[52]
Napoléon	332 min (5 hr, 32 min)	1927 (Restoration, 2016)	[53]
1900	317 min (5 hr, 17 min)	1976	[54]
Happy Hour	317 min (5 hr, 17 min)	2015	[55]
Batang West Side	315 min (5 hr, 15 min)	2001	[56]
The Deluge	315 min (5 hr, 15 min)	1974	[57]
Fanny and Alexander	312 min (5 hr, 12 min)	1982	[58]
Tsahal	304 min (5 hr, 4 min)	1994	[59]

Beyond List Pages

- Many pages contain list-like constructs
- Usually
  - small
  - same type
  - same relation to page entity
  - more grey links

Axl Rose

From Wikipedia, the free encyclopedia

Discography [edit]

with Guns N' Roses [edit]

- *Appetite for Destruction* (1987)
- *G N' R Lies* (1988)
- *Use Your Illusion I* (1991)
- *Use Your Illusion II* (1991)
- *"The Spaghetti Incident?"* (1993)
- *Chinese Democracy* (2008)

with Hollywood Rose [edit]

- *The Roots of Guns N' Roses* (2004)

with Rapidfire [edit]

- *Ready to Rumble* EP (2014)

Guest appearances [edit]

- *The Decline of Western Civilization Part II: The Metal Years – Original Motion Picture Soundtrack* by various artists (1988; "Under My Wheels" ft. Alice Cooper, Slash and Izzy Stradlin)
- *The End of the Innocence* by Don Henley (1989; "I Will Not Go Quietly")
- *Fire and Gasoline* by Steve Jones (1989; "I Did U No Wrong")
- *Pawnshop Guitars* by Gilby Clarke (1994; "Dead Flowers")
- *Anxious Disease* by The Outpatience (1996; "Anxious Disease" ft. Slash)
- *Angel Down* by Sebastian Bach (2007; "Back in the Saddle," "(Love Is) a Bitchslap," "Stuck Inside")
- *New Looney Tunes* (2018, "Rock the Rock")[122]

Filmography [edit]

Title	Year	Role	Notes
The Dead Pool	1988	Musician at funeral	Uncredited
Grand Theft Auto: San Andreas (video game)	2004	DJ Tommy "The Nightmare" Smith in the K-DST radio	Voice
That Metal Show	2011	Himself	
Jimmy Kimmel Live!	2012	Himself	
New Looney Tunes (TV show)[123]	2018	Himself	Voice
Scooby-Doo and Guess Who? (Tv Show)	2021	Himself	Voice

Beyond List Pages

(artist, Axl_Rose)

$\exists topSection.\{„Discography“\} \sqsubseteq \exists artist.\{>PageEntity<\}$

(type, MusicalWork)

$\exists topSection.\{„Discography“\} \sqsubseteq MusicalWork$

(musicalBand, Guns_N'_Roses)

$\exists topSection.\{„Discography“\} \sqcap sectionEntityType.\{Band\} \sqsubseteq \exists musicalBand.\{>SectionEntity<\}$

Axl Rose — Page Entity

From Wikipedia, the free encyclopedia

Discography [edit] — Top Section

with Guns N' Roses [edit] — Section

- *Appetite for Destruction* (1987)
- *G N' R Lies* (1988)
- *Use Your Illusion I* (1991)
- *Use Your Illusion II* (1991)
- *"The Spaghetti Incident?"* (1993)
- *Chinese Democracy* (2008)

with Hollywood Rose [edit]

- *The Roots of Guns N' Roses* (2004)

with Rapidfire [edit]

- *Ready to Rumble* EP (2014)

Guest appearances [edit]

- *The Decline of Western Civilization Part II: The Metal Years – Original Motion Picture Soundtrack* by various artists (1988; "Under My Wheels" ft. Alice Cooper, Slash and Izzy Stradlin)
- *The End of the Innocence* by Don Henley (1989; "I Will Not Go Quietly")
- *Fire and Gasoline* by Steve Jones (1989; "I Did U No Wrong")
- *Pawnshop Guitars* by Gilby Clarke (1994; "Dead Flowers")
- *Anxious Disease* by The Outpatience (1996; "Anxious Disease" ft. Slash)
- *Angel Down* by Sebastian Bach (2007; "Back in the Saddle," "(Love Is) a Bitchslap," "Stuck Inside")
- *New Looney Tunes* (2018, "Rock the Rock")[122]

Filmography [edit]

Title	Year	Role	Notes
The Dead Pool	1988	Musician at funeral	Uncredited
Grand Theft Auto: San Andreas (video game)	2004	DJ Tommy "The Nightmare" Smith in the K-DST radio	Voice
That Metal Show	2011	Himself	
Jimmy Kimmel Live!	2012	Himself	
New Looney Tunes (TV show)[123]	2018	Himself	Voice
Scooby-Doo and Guess Who? (Tv Show)	2021	Himself	Voice

Beyond List Pages

- Learning descriptive rules for listings, e.g.
  - topSection(“Discography”) → artist.{PageEntity}
  - Learning across pages to mitigate small data problems
- Metrics:
  - Support: no. of listings covered by rule antecedent
  - Confidence: frequency of rule consequent over all covered listings
  - Consistency: mean absolute deviation of overall confidence and listing confidence
    - i.e., does the rule work equally well across all covered listings

CaLiGraph at a Glance

- Latest version 2.1
  - 15M entities
    - incl. 8M from listings
  - Caveat:
    - disambiguation!

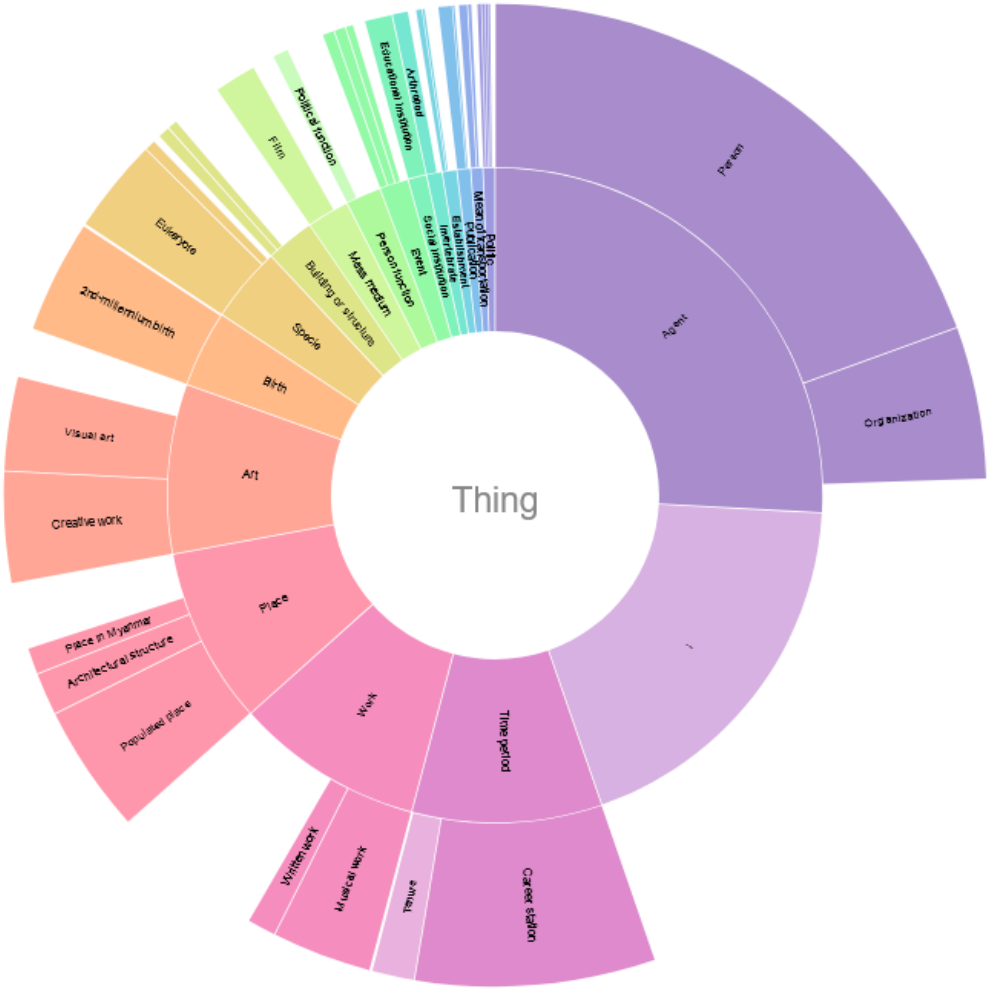

Entity Disambiguation

- Examples: Wikipedia pages of *Die Krupps* and *Eisbrecher*

Members [edit]

- Jürgen Engler – vocals, guitar, keyboards, synthesizers and programming, metallic percussion (1980–1985, 1989–1997, 2005–present)
- Ralf Dörper – keyboards, synthesizers and programming (1980–1982, 1985, 1989–1997, 2005–present)
- Marcel Zürcher – guitar, keyboards (2005–present)
- Nils Finkeisen – guitar (2015–present)
- Paul Keller – drums (2018–present)

Former members [edit]

- Bradley Bills – live drums (2013–2014)
- Rüdiger Esch – bass guitar (1989–1997, 2005–2011)
- Christoph "Nook" Michelfeit – drums, electronic percussion
- Bernward Malaka – bass guitar (1980–1982)
- Hendrik Thiesbrummel – live drums (2016–2018)
- Frank Köllges – drums
- Eva Gossling – saxophone (1981)
- Tina Schnekenburger – syncussion, bass
- Walter Jäger – ?
- Christopher Lietz – programming, samples (1995–1997)
- Lee Altus – guitar (1992–1997)
- Darren Minter – drums (1993)
- George Lewis – drums (1997)
- Oliver Röhl – drums
- Achim Färber – drums
- Volker Borchert – drums (1992, 2015–2016)

?

Members [edit]

- Alexx Wesselsky – vocals (2003–present)
- Noel Pix – lead guitar, programming, production (2003–present)
- Jürgen Plangger – rhythm guitar (2007–present)
- Maximilian Schauer – keyboards, programming (live and session: 2003–2007, session only: 2008–present)
- Achim Färber – drums (2011–present)
- Rupert Keplinger – bass (2013–present)

Eisbrecher, line-up at Rockharz Open Air 2018


Alexander "Alexx" Wesselsky


Jochen "Noel Pix" Seibert


Jürgen Plangger


Achim Färber

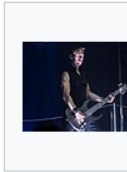
Rupert Keplinger

Former live members [edit]

- Felix Primc – rhythm guitar (2003–2007)
- Micheal Behnke – bass (2003–2007)
- Martin Motnik – bass (2007–2008)
- Olli Pohl – bass (2008–2010, 2015)
- Dominik Palmer – bass (2010–2013)
- Rene Greil – drums (2003–2011)

Touring members [edit]

- Sebastien Angrand – drums (2010)

CaLiGraph Glitches

CaLiGraph

Formats

Sparql Endpoint

About: clgr:Mannheim

Property	Value
rdfs:label	• Mannheim
clgo:country	• Moldova • Germany
rdf:type	• Planned capital • City in Baden-Würt • Twin town or sister • Coat of arms with the Palatine Lion • French exonym for German toponyms • Twin town or sister city in Lithuania • University town in Germany • owl:NamedIndividual • City or town in Germany • Most polluted city in the world

List of twin towns and sister cities in Moldova

From Wikipedia, the free encyclopedia

This is a list of places in Moldova having standing links to local communities in other countries. In most cases, the association, especially when formalised by lo government, is known as "town twinning" (though other terms, such as "partner towns" or "sister cities" are sometimes used instead), and while most of the pla are towns, the list also comprises villages, cities, districts, counties, etc. with similar links.

Index: A · B · C · D · E · F · G · H · I · J · K · L · M · N · O · P · Q · R · S · T · U · V · W · X · Y · Z · References

C [edit]

Chișinău[3][4]

- Alba Iulia, Romania
- Ankara, Turkey
- Bucharest, Romania
- Chernivtsi, Ukraine
- Grenoble, France
- Iași, Romania
- Kiev, Ukraine
- Mannheim, Germany
- Minsk, Belarus
- Odessa, Ukraine
- Reggio Emilia, Italy
- Sacramento, United States
- Tbilisi, Georgia
- Tel Aviv, Israel
- Vilnius, Lithuania
- Yerevan, Armenia

Comrat

- Bălți, Moldova[1]

Criuleni

- Jurbarkas, Lithuania[5]

From DBpedia to DBkWik

- Wikipedia-based Knowledge Graphs will remain an essential building block of Semantic Web applications
- But they suffer from...
  - ...a coverage bias
  - ...limitations of the creating heuristics

From DBpedia to DBkWik

- One (but not the only!) possible source of coverage bias
  - Articles about long-tail entities become deleted

Wikipedia:Deletion policy

From Wikipedia, the free encyclopedia

"WP:DP", "WP:DELETE", and "WP:DEL" redirect here. For other uses, see Wikipedia:Deletion pr Delaware.

This page documents an English Wikipedia policy.
It describes a widely accepted standard that all editors should normally

This page in a nutshell: Administrators have the ability to delete in accordance with established policies and guidelines, and comm

The **Wikipedia deletion policy** describes how articles, media, and other pages that do not meet the each day through the processes outlined below.

Deletion of a Wikipedia article removes the current version and all pr reverse ("undelete") any deletion. All such actions (other than viewing administrators normally will not delete it.

Contents [hide]

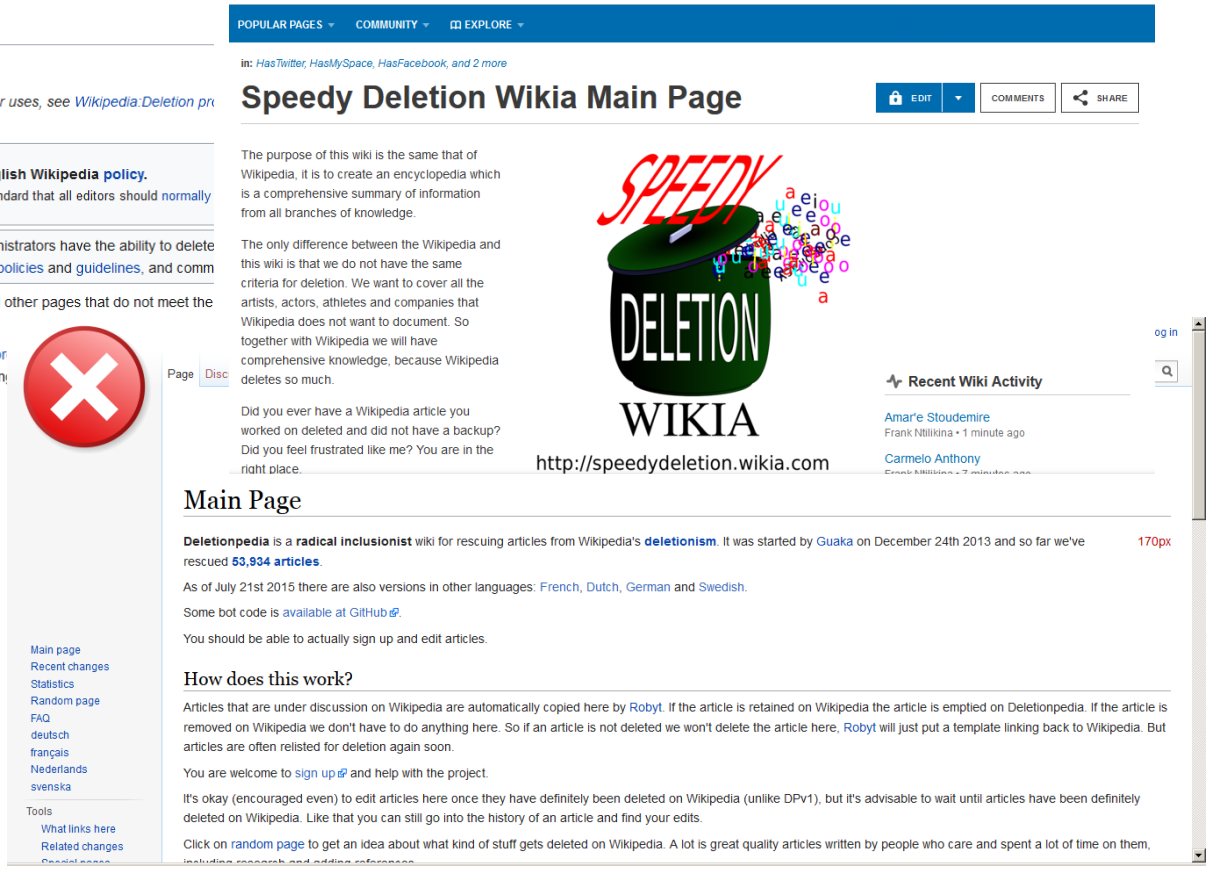

POPULAR PAGES COMMUNITY EXPLORE

in: *HasTwitter, HasMySpace, HasFacebook, and 2 more*

Speedy Deletion Wikia Main Page

EDIT COMMENTS SHARE

The purpose of this wiki is the same that of Wikipedia, it is to create an encyclopedia which is a comprehensive summary of information from all branches of knowledge.

The only difference between the Wikipedia and this wiki is that we do not have the same criteria for deletion. We want to cover all the artists, actors, athletes and companies that Wikipedia does not want to document. So together with Wikipedia we will have comprehensive knowledge, because Wikipedia deletes so much.

Did you ever have a Wikipedia article you worked on deleted and did not have a backup? Did you feel frustrated like me? You are in the right place.

SPEEDY DELETION WIKIA

http://speedydeletion.wikia.com

Recent Wiki Activity

Amar'e Stoudemire
Frank Ntilikina • 1 minute ago

Carmelo Anthony

Page Disc

Main page
Recent changes
Statistics
Random page
FAQ
deutsch
français
Nederlands
svenska

Tools
What links here
Related changes

Main Page

Deletionpedia is a **radical inclusionist** wiki for rescuing articles from Wikipedia's **deletionism**. It was started by Guaka on December 24th 2013 and so far we've rescued **53,934 articles**.

As of July 21st 2015 there are also versions in other languages: French, Dutch, German and Swedish.

Some bot code is available at GitHub.

You should be able to actually sign up and edit articles.

How does this work?

Articles that are under discussion on Wikipedia are automatically copied here by Robyt. If the article is retained on Wikipedia the article is emptied on Deletionpedia. If the article is removed on Wikipedia we don't have to do anything here. So if an article is not deleted we won't delete the article here, Robyt will just put a template linking back to Wikipedia. But articles are often relisted for deletion again soon.

You are welcome to sign up and help with the project.

It's okay (encouraged even) to edit articles here once they have definitely been deleted on Wikipedia (unlike DPv1), but it's advisable to wait until articles have been definitely deleted on Wikipedia. Like that you can still go into the history of an article and find your edits.

Click on random page to get an idea about what kind of stuff gets deleted on Wikipedia. A lot is great quality articles written by people who care and spent a lot of time on them,

From DBpedia to DBkWik

- Why stop at Wikipedia?
- Wikipedia is based on the MediaWiki software
  - ...and so are thousands of Wikis
  - Fandom by Wikia: >385,000 Wikis on special topics
  - WikiApiary: reports >20,000 installations of MediaWiki on the Web

From DBpedia to DBkWik

- Collecting Data from a Multitude of Wikis

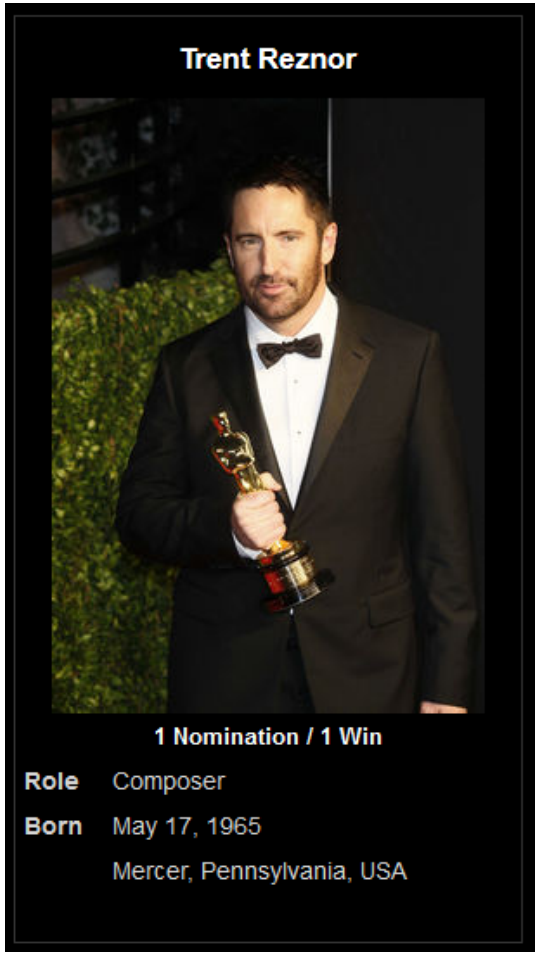

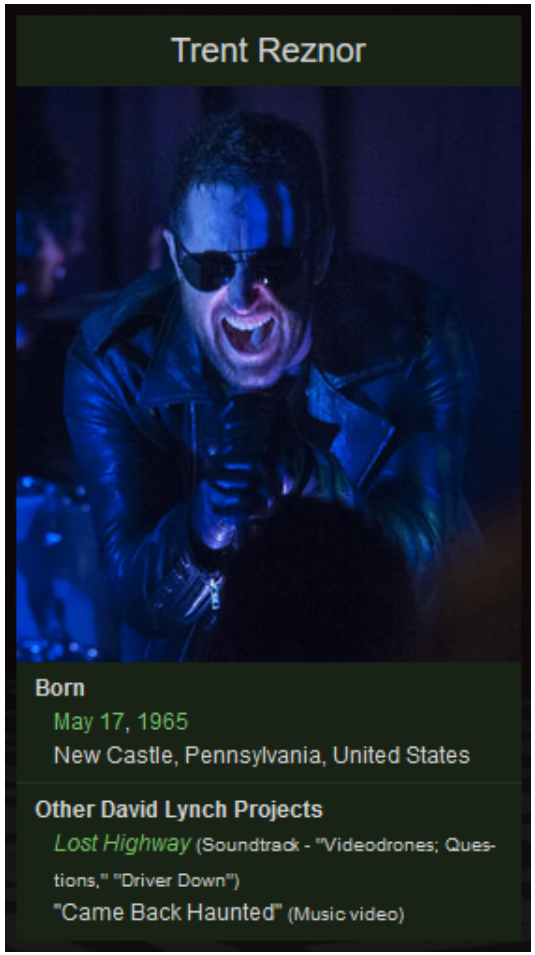

From DBpedia to DBkWik

- The DBpedia Extraction Framework consumes MediaWiki dumps
- Experiment (started as team project 2017)
  - Can we process dumps from arbitrary Wikis with it?
  - Are the results somewhat meaningful?

From DBpedia to DBkWik

- Example from Harry Potter Wiki

DBkWik Browse using Formats Faceted Browser Sparql Endpoint

About: http://dbkwik.webdatacommons.org/HarryPotter/resource/Gryffindor

Property	Value
owl:sameAs	• dbr:Gryffindor
foaf:depiction	• http://commons.wikimedia.org/wiki/Special:FilePath/0.31_Gryffindor_Crest_Transparent.png
dcterms:subject	• dbkwik:HarryPotter/resource/Category:Gryffindor_House • dbkwik:HarryPotter/resource/Category:Hogwarts_Houses
skos:altLabel	• Gryffindor House • Gryffindor Students • Griffindor • Gryffindors • Griffyndor • Griffyndor Students • Gryffindor Girl • The Gryffindors
skos:prefLabel	• Gryffindor
dbkwik:HarryPotter/ontology/thumbnail	• http://commons.wikimedia.org/wiki/Special:FilePath/0.31_Gryffindor_Crest_Transparent.png?width=300
dbkwik:HarryPotter/property/animal	• dbkwik:HarryPotter/resource/Lion

http://dbkwik.org/

From DBpedia to DBkWik

- Differences to DBpedia
  - DBpedia has manually created mappings to an ontology
  - Wikipedia has one page per subject
  - Wikipedia has global infobox conventions (more or less)

- Challenges
  - On-the-fly ontology creation
  - Instance matching
  - Schema matching

Hertling & Paulheim: *DBkWik: A Consolidated Knowledge Graph from Thousands of Wikis.* ICBK 2018

From DBpedia to DBkWik

- Heuristics
  - Ontology induction
  - Instance/Schema Matching

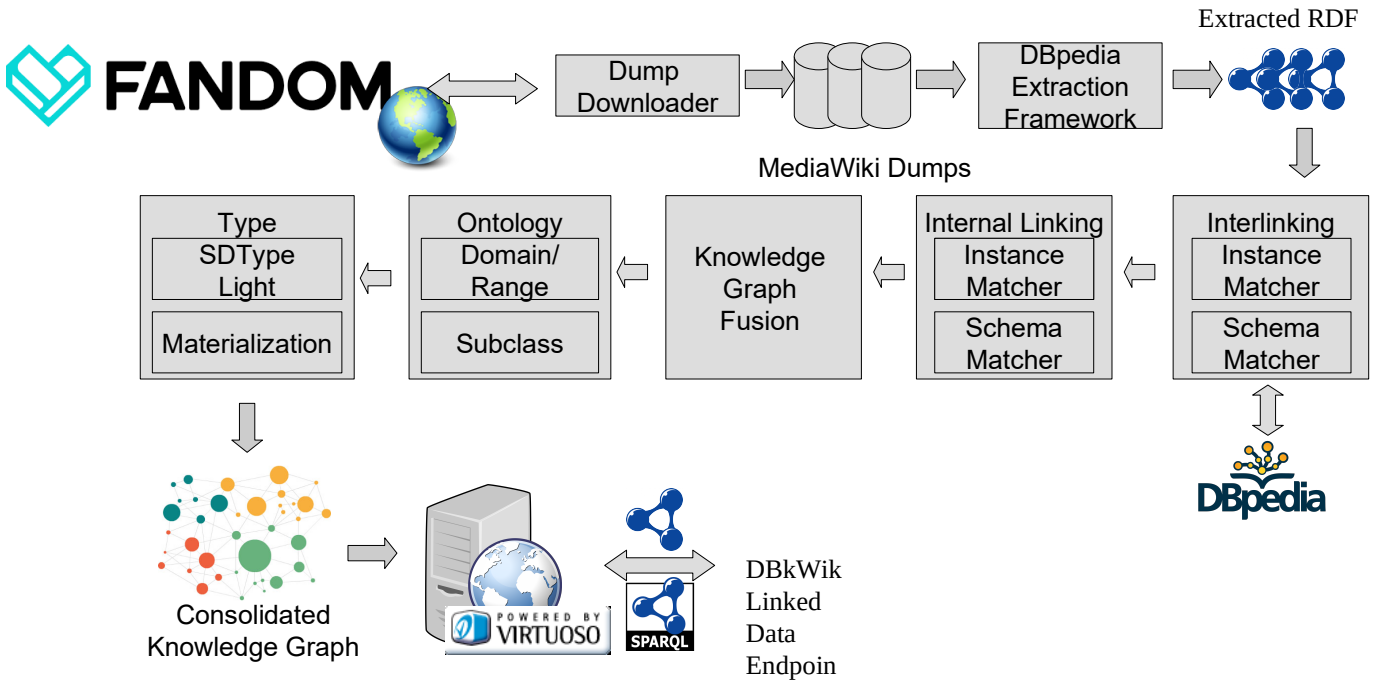

Hertling & Paulheim: *DBkWik: A Consolidated Knowledge Graph from Thousands of Wikis.* ICBK 2018

From DBpedia to DBkWik

- Downloaded ~15k Wiki dumps from Fandom
  - 52.4GB of data, roughly the size of the English Wikipedia

- Prototype: extracted data for ~250 Wikis
  - 4.3M instances, ~750k linked to DBpedia
  - 7k classes, ~1k linked to DBpedia
  - 43k properties, ~20k linked to DBpedia
  - ...including duplicates!

- Link quality
  - Good for classes, OK for properties (F1 of .957 and .852)
  - Needs improvement for instances (F1 of .641)

From DBpedia to DBkWik

- Scalability of matching:
  - Pairwise matching does not scale
  - 300k Wikis, 1 minute for each pair → 171k years
- Iteratively match and merge
  - 300k Wikis, 1 minute for each match&merge run → 200 days
- Tree-shaped execution plan
  - Parallelizable
  - Hierarchical clustering by topic
  - Whole run under a week

WebIsALOD

- Background: Web table interpretation
- Most approaches need typing information
  - DBpedia etc. have too little coverage on the long tail
  - Wanted: extensive type database

Rank	Country/Territory	Capital	Population	Year	Percent of Population
1	China	Beijing	20,693,000[1]	2012	1.52%
2	India	New Delhi	16,787,949[2]	2014	0.90%
3	Japan	Tokyo	13,189,000[3]	2011	10.32%
4	Philippines	Manila	12,877,253[4]	2015	12.44%
5	Russia	Moscow	11,541,000[5]	2011	8.07%
6	Egypt	Cairo	10,230,350	2012	11.10%
7	Indonesia	Jakarta	10,187,595[6]	2011	4.18%
8	Democratic Republic of the Congo	Kinshasa	10,125,000[7]	2012	12.30%
9	South Korea	Seoul	9,989,795[8]	2015	20.47%
10	Bangladesh	Dhaka	8,906,000 [9]	2011	5.56%
11	Mexico	Mexico City	8,851,080[10]	2010	7.51%
12	Iran	Tehran	8,846,782	2014	9.91%
13	United Kingdom	London	8,630,100[11]	2015	13.25%
14	Peru	Lima	8,481,415[12]	2012	28.29%
15	Thailand	Bangkok	8,249,117[13]	2010	12.42%
16	Colombia	Bogotá	7,613,303[14]	2011	16.17%
17	Vietnam	Hanoi	7,587,800[15]	2014	8.22%
18	*Hong Kong (China)*	Hong Kong	7,298,600[16]	2015	100%
19	Iraq	Baghdad	7,216,040[17]		21.59%
20	Singapore	Singapore	5,535,000[18]	2015	100%
21	Turkey	Ankara	5,150,072	2014	6.72%
22	Chile	Santiago	5,084,038[19]	2012	29.12%
23	Saudi Arabia	Riyadh	4,878,723[20]	2009	18.20%
24	Germany	Berlin	3,520,000[21]	2012	4.38%
25	Syria	Damascus	3,500,000		15.32%
26	Algeria	Algiers	3,415,811		8.45%
27	Spain	Madrid	3,233,527[22]	2012	6.84%
28	North Korea	Pyongyang	3,144,005		12.63%
29	Afghanistan	Kabul	3,140,853		10.28%
30	Kenya	Nairobi	3,138,369	2010	7.67%

Hertling & Paulheim: *WebIsALOD: Providing Hypernymy Relations extracted from the Web as Linked Open Data.* ISWC 2017

WebIsALOD

- Extraction of type information using Hearst-like patterns, e.g.,
  - T, such as X
  - X, Y, and other T
- Text corpus: common crawl
  - ~2 TB crawled web pages
  - Fast implementation: regex over text
  - “Expensive” operations only applied once regex has fired
- Resulting database
  - 400M hypernymy relations

Common Crawl

Seitner et al.: *A large DataBase of hypernymy relations extracted from the Web.* LREC 2016

WebIsALOD

- Example:

About: fiction writer

Premodifier: fiction
Head noun: writer

Same concepts
http://dbpedia.org/resource/Fiction

Broader concepts

label	provenance	confidence
writer	isap:391280092	0.799331
great idea	isap:493244047	0.672180
several magazine	isap:104101164	0.655684
category	isap:104762336	0.477191
artist	isap:387107910	0.471280
blog	isap:492616562	0.458511
writers	isap:439522913	0.427701
story	isap:122402598	0.306667
group	isap:115379219	0.299656
poet	isap:492284397	0.287519

Narrower concepts

label	provenance	confidence
george orwell	isap:386468501	0.662121
science fiction	isap:275868279	0.635886
franz kafka	isap:159147340	0.602015
steve almond	isap:392552636	0.581515
dan brown	isap:157209267	0.574584
james joyce	isap:159394667	0.561794
stephen king	isap:306753456	0.557354
flannery oconnor	isap:266705231	0.555096
alice munro	isap:162537618	0.552608
ayn rand	isap:301402665	0.526857

http://webisa.webdatacommons.org/

WebIsALOD

- Initial effort: transformation to a LOD dataset
  - including rich provenance information

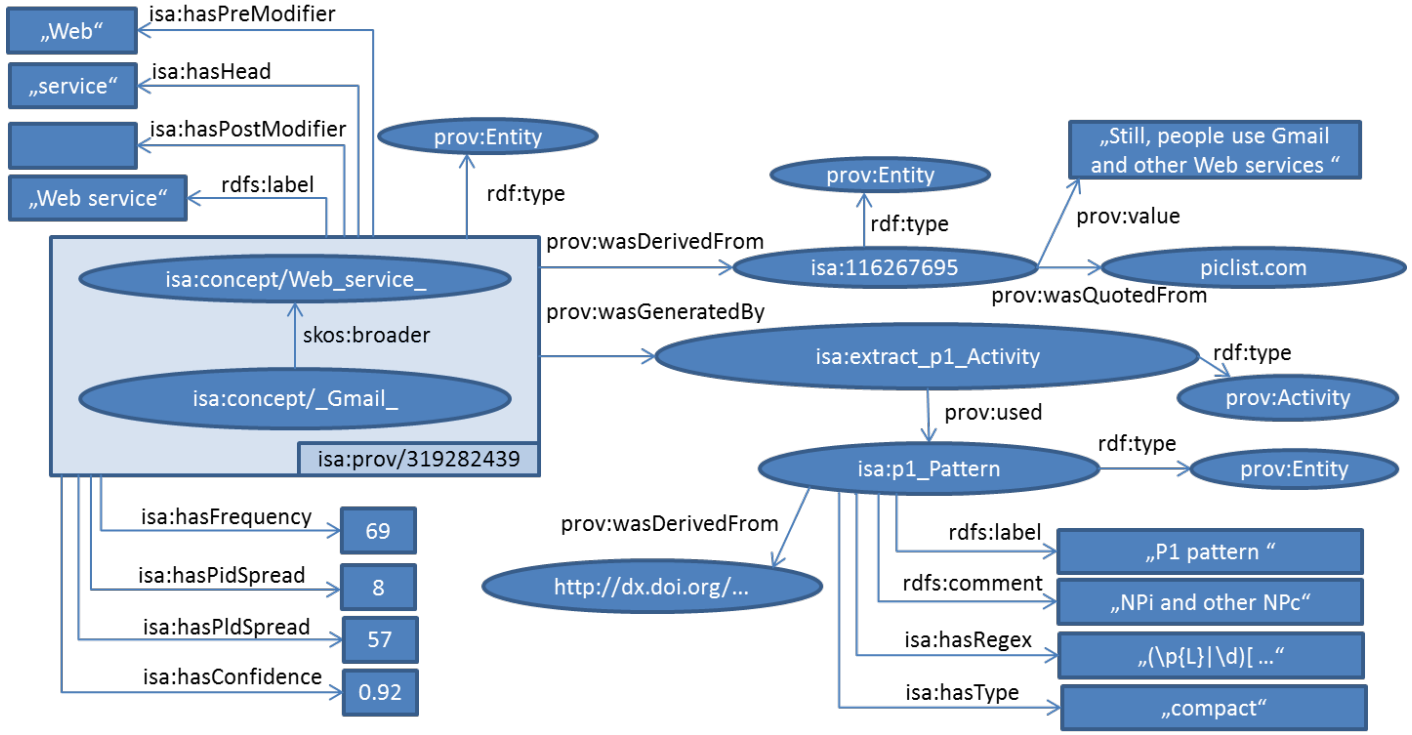

Hertling & Paulheim: *WebIsALOD: Providing Hypernymy Relations extracted from the Web as Linked Open Data.* ISWC 2017

WebIsALOD

- Estimated contents breakdown

Hertling & Paulheim: *WebIsALOD: Providing Hypernymy Relations extracted from the Web as Linked Open Data.* ISWC 2017

WebIsALOD

- Main challenge
  - Original dataset is quite noisy (<10% correct statements)
  - Recap: coverage vs. accuracy
  - Simple thresholding removes too much knowledge
- Approach
  - Train RandomForest model for predicting correct vs. wrong statements
  - Using all the provenance information we have
  - Use model to compute confidence scores
- Current ongoing research
  - Using transformers and a larger training set

Hertling & Paulheim: *WebIsALOD: Providing Hypernymy Relations extracted from the Web as Linked Open Data.* ISWC 2017

WebIsALOD

- Current challenges and works in progress
  - Distinguishing instances and classes
    - i.e.: subclass vs. instance of relations
  - Splitting instances
    - *Bauhaus is a goth band*
    - *Bauhaus is a German school*
  - Knowledge extraction from pre and post modifiers
    - *Bauhaus is a goth band* → genre(Bauhaus, Goth)
    - *Bauhaus is a German school* → location(Bauhaus, Germany)

Hertling & Paulheim: *WebIsALOD: Providing Hypernymy Relations extracted from the Web as Linked Open Data.* ISWC 2017

Summary

- We have seen a couple of Knowledge Graphs
  - How they are built
  - What they contain
- For your project
  - Have a look at the fit for your domain
  - Try different options
- For a master's thesis later
  - Work on recent developments in our group

Questions?