Public Knowledge Graphs



IE650 Knowledge Graphs



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



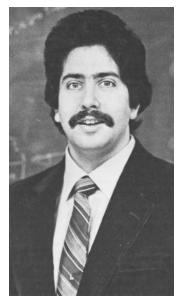
• Knowledge Graphs out there (not guaranteed to be complete)

Name	Instances	Facts	Types	Relations	
DBpedia (English)	4,806,150	176,043,129	735	2,813	7
YAGO	4,595,906	25,946,870	488,469	77	
Freebase	49,947,845	3,041,722,635	26,507	37,781	– public
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	– private
Yahoo! Knowledge Graph	3,443,743	1,391,054,990	250	800	

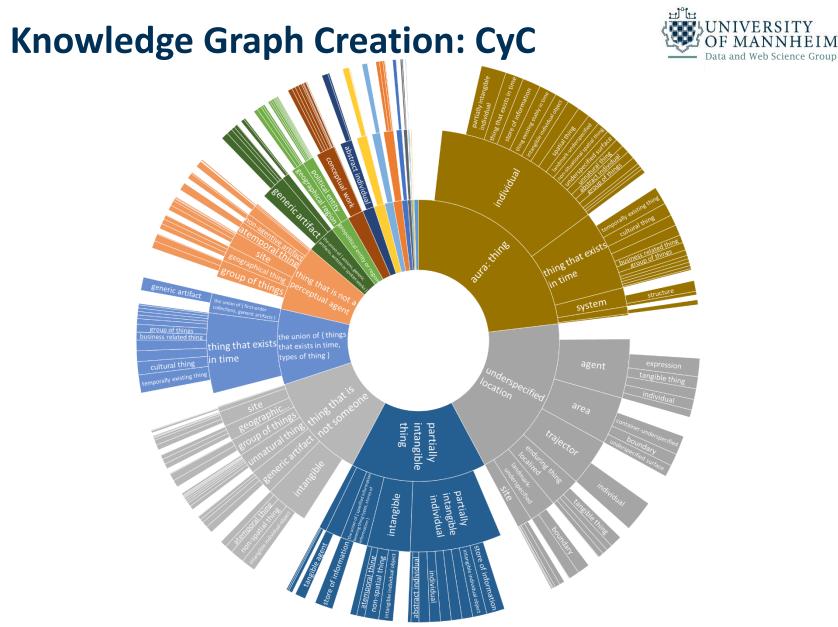
University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.20243Paulheim: Knowledge graph refinement: A survey of approaches and evaluation methods. Semantic Web 8:3 (2017), pp. 489-508

Knowledge Graph Creation: CyC



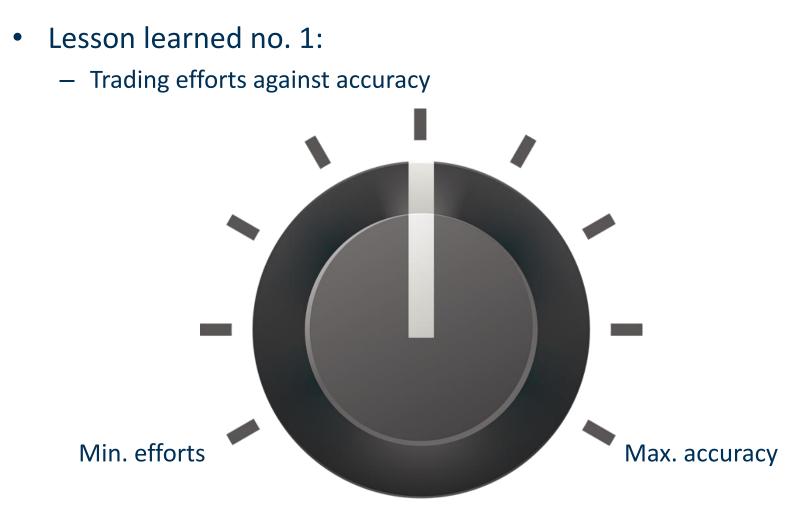


- The beginning
 - Encyclopedic collection of knowledge
 - Started by Douglas Lenat in 1984
 - Estimation: 350 person years and 250,000 rules should do the job of collecting the essence of the world's knowledge
- The present (as of June 2017)
 - ~1,000 person years, \$120M total development cost
 - 21M axioms and rules
 - Used to exist until 2017



Knowledge Graph Creation





Knowledge Graph Creation: Freebase

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



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

Freebase

Knowledge Graph Creation: Freebase



- Community based
- Like Wikipedia, but more structured

Arnold Schwa	arzenedder -
🔍 Discuss "Arnold Schv	varzenegger" 🗉 Show Empty Fields
	 = 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
To Kat	.≡ Date of Birth: Jul 30, 1947
THE REAL PROPERTY AND	.≡ Place of Birth: Thal, Austria
	.= Country Of Nationality: United States
< image 1 of 1 🕨	.= 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



Knowledge Graph Creation: Wikidata



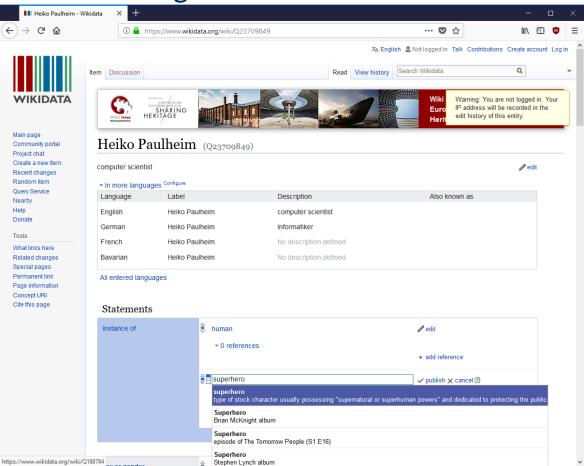
- The 2010s
 - Wikidata: launched 2012
 - Goal: centralize data from Wikipedia languages
 - Collaborative
 - Imports other datasets
- Present
 - One of the largest public knowledge graphs (see later)
 - Includes rich provenance



Knowledge Graph Creation: Wikidata



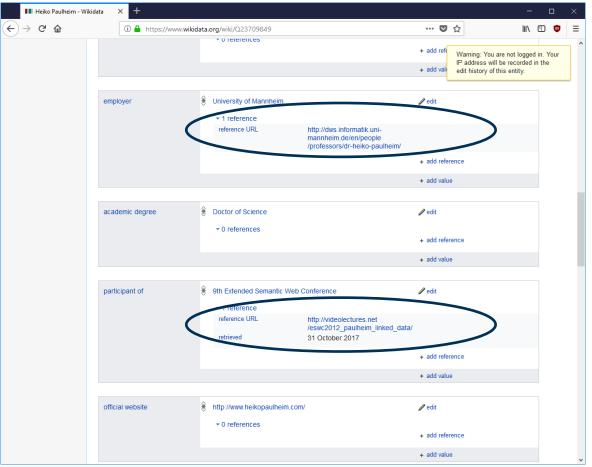
• Collaborative editing

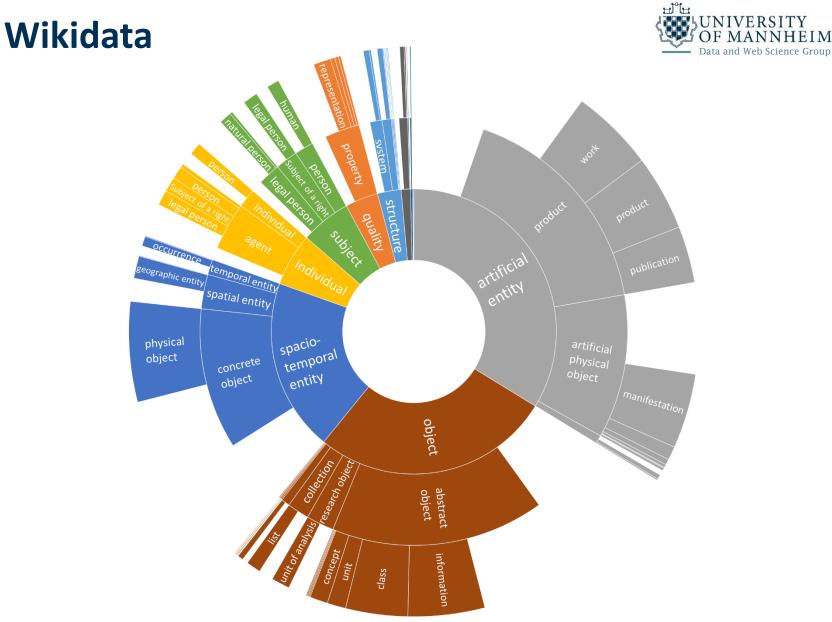


Knowledge Graph Creation: Wikidata



• Provenance





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







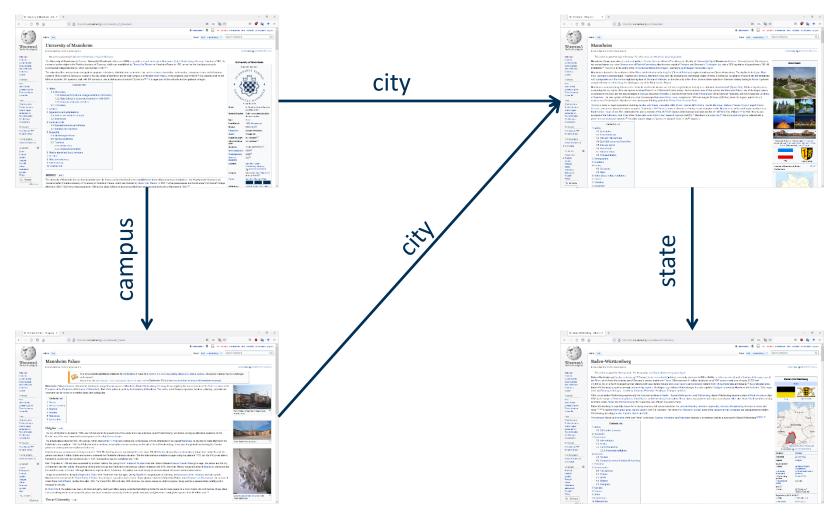
Wikipedia as a Knowledge Graph



W University of Ma	annheim - Wikip × +				-	
\rightarrow C D	C A https://en.wikipedia.org/wiki/Unive	rsity_of_Mannheim 🗄 90% 🖏	23	▽ 🤘))	•
9 . 01 0		🚨 HeikoPaulheim 🌲 🗔 Taik Sar	ndbox Preferences Bet	ta Watchlist C	ontribution	s Log (
A W S	Article Talk	Read Edit View history 🛱	Search Wikipedia			Q
a an 2						
VIKIPEDIA	University of Mannheim					
e Free Encyclopedia	From Wikipedia, the free encyclopedia		Cor	ordinates: 🔔 49	29'00"N 8	*27'53*/
n page itents	Not to be confused with Mannheim University of Applied Sciences					
rrent events		abbreviated UMA, is a public research university in Mannheim, Baden-Württemberg, Germany. Founded in 1967, the	e Univer	sity of Man	nheim	
ndom article		was established by Elector Carl Theodor at Mannheim Palace in 1763, as well as the Handelshochschule	Uni	iversität Mannhe	im	
out Wikipedia	(Commercial College Mannheim), which was founded in 1907. ^[4]			AIBUS V		
ntact us nate	, , , , , , , , , , , , , , , , , , , ,	is administration, economics, law, social sciences, humanities, mathematics, computer science and information	OMA	V	RI	
		Inheim and its main campus is in the Mannheim Palace. In the academic year 2016/2017 the university had 12,000	, Î	⊷ ^{t⁰t} ⊷	T.	
tribute	full-time students, 907 academic staff, with 194 professors, and a tota	al income of around €123 million. ^{[2][1]} It is organized into five schools and two graduate colleges.				
) 	Contents [hide]		<u>الم</u> •		4 •	
rn to edit nmunity portal	1 History		4			
ent changes	1.1 20th century		UZ E			
bad file	1.1.1 Municipal Commercial College Mannheim (1907–1933)		14	TĂT MN		
Is	1.1.2 State College for Economics Mannheim (1946–1967)		CRS.	ITT NON	44	
at links here	1.1.3 University of Mannheim (1967)			Seal of the UMA		
ated changes	1.2 21st century		Motto	In Omnibus \	Verites Su	orema
cial pages	2 Campus			Lex Esto (Lat		
manent link	3 Organisation and administration		Motto in English	Truth in every	thing shou	uld be
e information	3.1 Schools and Graduate Colleges			the supreme I	aw	
this page data item	3.2 Governance		Туре	Public		
idata item	4 Academic profile 4.1 Research institutes and affiliates		Established	1967; 54 year	rs ago	
t/export	4.2 Rankings and reputation		Budget	€123 million ^[1]	1	
vnload as PDF	5 Student life		Chancellor	Barbara Wind	scheid	
table version	5.1 Student organizations		Rector	Thomas Puhl		
ther projects	5.2 Sports and athletics		Academic staff	907 (full time)		
imedia Commons	5.3 Traditions		Administrative staff	617 (full time)	[1]	
	5.3.1 Schlossfest		Students	12,000 (HWS 2	2020/240[2]	
guages 🗘	5.3.2 Schneckenhof Parties		Undergraduates			
sk	6 Notable alumni and faculty members		Postgraduates	4.828[2]		
tsch añol	7 See also		Doctoral	793[2]		
nçais	8 Notes and references		students			
atski	9 Further reading		Location	Mannheim, Ba		
ano	10 External links			Württemberg,		
derlands сский				(a) 49°29'00*		
скии kce	History [edit]		Campus	Urban (Mann) acres (0.3 km		ð), 74
·		ne he deted back to the establishment of one of its reader one institutions when <i>Hermiter</i> to Alexandratic dec	Colors	Mannheim Blu		te
19 more		can be dated back to the establishment of one of its predecessor institutions - the Kurpfälzische Akademie der				
Edit links	wissenscharten (Palatine Academy of Sciences) in Mannheim Palac	e, which was founded by Elector Carl Theodor in 1763. Further predecessors are the Municipal Commercial Colleg	3			

Wikipedia as a Knowledge Graph

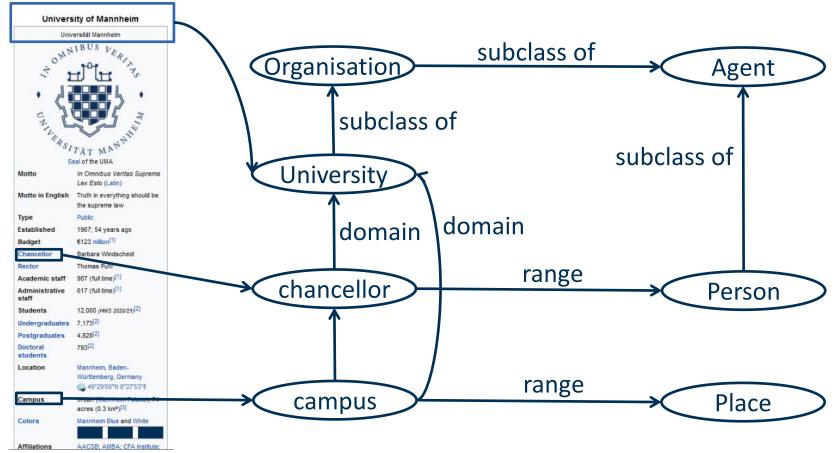




Wikipedia as a Knowledge Graph



Mapping to a central schema/ontology



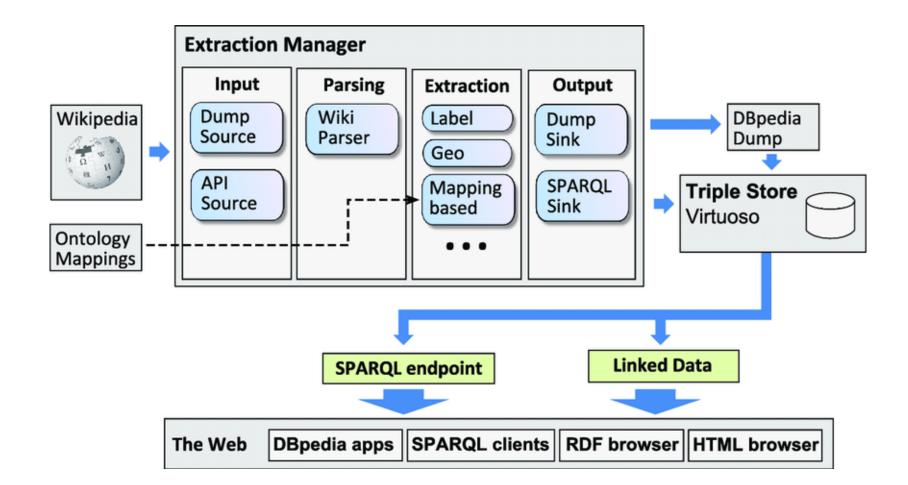


Univers	ity of Mannheim	{{Infobox univer	rsity					
	-	motto	=''In Omnibus Veritas Suprema Lex Esto'' ([[Latin]])					
Unive	ersität Mannheim	mottoeng	= Truth in everything should be the supreme law					
14 M	DUS VEN	name	=University of Mannheim					
ONN		native name	=Universität Mannheim					
∠ <		image name	=Uni Mannheim Siegel.gif					
	s s s s s s s s s s s s s s s s s s s	caption	=[[Seal (emblem) Seal]] of the UMA					
		established	=1763: Theodoro Palatinae 1907: Handelshochschu					
z		type	=[[Public University Public]]					
-2 -2		endowment	E115 [[million]]					
	2 2	academic staff	=800 (full time)					
ERS .	× 4	administrative	<pre>staff = 550 (full time)</pre>					
1	AT M Carlor	Schools						
		rector	=[[Ernst-Ludwig von Thadden]]					
Motto	In Omnibus Veritas Suprema Lex Esto (Latin)	chancellor	=[[Susann-Annette Storm]]					
Motto in English	Truth in everything should	students	=12,151 <small>''(HWS 2013/14)''</small> <ref name="ur</th></tr><tr><th>motto in English</th><th>be the supreme law</th><th>/Studierendensta</th><th>atistik_hws13.pdf title= Studierendenstatistik der Univ</th></tr><tr><th>Established</th><th>1763: Theodoro Palatinae</th><th>undergrad</th><th>=6,915<ref name=" uni-mannheim.de"=""></ref>					
	1907: Handelshochschule	postgrad	=6,915 <ref name="un1-mainneim.de"></ref> =4,965 <re; <2xml="" ?="" encoding="utf-8" version="1.0"></re;>					
	1967: Universität	doctoral	=249 <ref 1="" <rdf:description="" rdf:about="http://dbpedia.org/resource/University_of_Mannheim"></ref>					
	Mannheim	profess	<pre>= <dbp:mottoeng xml:lang="en">Truth in everything should be the supreme law</dbp:mottoeng></pre>					
Туре	Public	city	=[[Mannhe: <dbp:chancellor xml:lang="en">Susann-Annette Storm</dbp:chancellor>					
Endowment	€115 million	state	<pre><dbp:budget rdf:datatype="http://dbpedia.org/datatype/euro">1.23E8</dbp:budget> <dbp:campus xml:lang="en">Urban , 74 acres</dbp:campus></pre>					
Chancellor	Susann-Annette Storm	country	=[[German] <dbp:caption xml:lang="en">Seal of the UMA</dbp:caption>					
Rector	Ernst-Ludwig von Thadden	coor	= {{Coord					
Academic staff	800 (full time)		<pre><dbp:colors xml:lang="en">Mannheim Blue and White</dbp:colors> <dbp:country rdf:resource="http://dbpedia.org/resource/Germany"></dbp:country></pre>					
Administrative	550 (full time)		<pre><dbp:country fd1:fesource="nttp://dbpedia.org/resource/dermany"></dbp:country> <dbp:imagename xml:lang="en">Uni Mannheim Seal.png</dbp:imagename></pre>					
staff			<pre><dbp:logo xml:lang="en">University of Mannheim.svg</dbp:logo></pre>					
Students	12,151 (HWS 2013/14) ^[1]		<pre><dbp:logosize rdf:datatype="http://www.w3.org/2001/XMLSchema#integer">280</dbp:logosize></pre>					
Undergraduates	6,915 ^[1]		<pre><dbp:nativename xml:lang="en">Universität Mannheim</dbp:nativename> <dbp:rector xml:lang="en">Thomas Puhl</dbp:rector></pre>					
Postgraduates	4,965 ^[1]		<pre><dbp:state rdf:resource="http://dbpedia.org/resource/Baden-Württemberg"></dbp:state></pre>					
Doctoral	249 ^[1]		<pre><dbp:type rdf:resource="http://dbpedia.org/resource/Public_University"></dbp:type></pre>					
students			<pre><dbp:website rdf:resource="https://www.uni-mannheim.de/en/%7C2=www.uni-mannheim.de"></dbp:website> </pre>					
			<pre><georss:point>49.4832 8.4647</georss:point></pre>					

</rdf:Description>

</rdf:RDF>





University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Lehmann et al.: DBpedia – A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia. 2014



Mapping en:Infobox film

OntologyProperty:director This is the mapping for the Wikipedia template Infobox film @. Find usages of this Wikipedia template Test this mapping & (or in namespace File & or Creator &) with some example Wikipedia pages. Che This is the definition of an ontology property. Read more about mapping Wikipedia templates.

Read more about editing the ontology schema.

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

Template Mapping (help)		Ontology	OntologyClass:Film		· · · · · · · · · · · · · · · · · · ·				
map to class	Film	OntologyC	Jass:Film	Ontology object property (help) rdfs:label (en) director					
		This is the deficitie			director				
Mappings				rdfs:label (en)	film director				
wappings		Show all properties	🚱 🖉 available for this class.	rdfs:label (nl)	regisseur				
Show class in class hierarchy 라. Read more about editing the ontology schema. You can see the result of your edit on DBpedia Live (rdfs:label (da)	instruktør						
		rdfs:label (de)	regisseur						
			rdfs:label (ru)	директор					
		rou can see the re	suit of your edit on DBpedia Live (rdfs:label (el)	σκηνοθέτης				
Property Mapping (h	elp)			rdfs:label (es)	director de cine				
		Ontology class (h		rdfs:label (fr)	réalisateur				
template property	director			rdfs:comment (en)	A film director is a person who directs the making of a film. ^[1]				
ontology property	director	rdfs:label (en)	film	rdfs:comment (fr)	Un réalisateur (au féminin, réalisatrice) est une personne qui dirige la fabrication d'une œuvre audiovisuelle, généralement pour le				
		rdfs:label (en)	movie	10 1	cinéma ou la télévision. ^[2] Film				
rdfs		rdfs:label (nl)	film	rdfs:domain					
		rdfs:label (da)	film	rdfs:range	Person				
		rdfs:label (de)	Film	rdf:type rdfs:subPropertyOf	dul:coparticipatesWith				
Property Mapping (h	elp)	rdfs:label (el)	ταινία	owl:equivalentProperty	schema:director, wikidata:P57				
template property	producer	rdfs:label (fr)	film	owl:propertyDisjointWit	h				
ontology property	producer	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

{{Infobox book

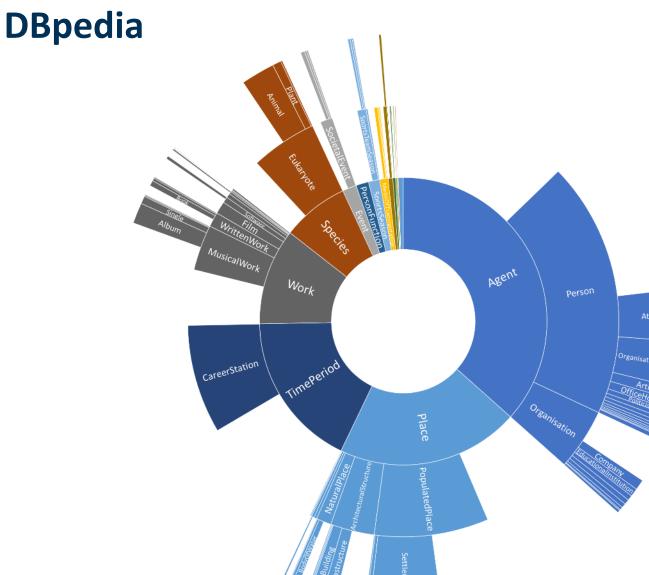
DBpedia 👁 Browse using 🔹 🖺 Formats 🔹

About: Tenth of December: Stories

An Entity of Type: Buch, from Named Graph: http://dbpedia.org, within Data Space: dbpedia.org

Tenth of December is a collection of short stories by American author George Sau magazines between 1995 and 2012. The book was published on January 8, 2013, 1 was a 2011 Bram Stoker Award finalist. Tenth of December was selected as one of New York Times Book Review. The collection also won the 2013 Story Prize for sho Folio Prize.

		{{Infobox book			was a 2011 brain stor	ter Award Infalist. Terrur of Decern	bel was selected as one of
Tenth o	of December	name = Tenth of [)ecember		New York Times Bool	k Review. The collection also won	the 2013 Story Prize for sho
	. 11.	author = [[George	e Saunders]]		Folio Prize.		
ler	nth	language = Englis	;h		-		
	I CI I	country = United	States		Property	Value	
of_	11071es	genre = [[Short s	tory collection]]		dbo:abstract		ction of short stories by America
De		publisher = [[Rar	ndom House]]				as published on January 8, 2013, cted as one of the 10 Best Books
	*30hrti5m, bizasar, and experience partets.	isbn = 0812993802	2				t-story collections and the inaug
		image = Tenth of	December.jpg				t story concetions and the mady
	0.0.0.0.0	caption = US rele	ase cover		dbo:author	dbr:George_Saunders	
	GEORGE SAUNDERS	release_date = Ja	anuary 8, 2013			-	
		media type = Prim	nt (hardcover)		<u>dbo:isbn</u>	• 0812993802	
am	hor <	pages = 208	>				
	iber 🗧	oclc = 785558855			<u>dbo:literaryGenre</u>	<u>dbr:Short_story_collection</u>	
	lease cover	Mapping en:In	fobox book		dbo:numberOfPages	 208 (xsd:positiveInteger) 	
Author	George Saunders					200 (Joaiposiareintegel)	
Language	English			ok⊠. Find usages of this Wikipedia templ		• 785558855	
Genre	Short story collection	Test this mapping der in m	amespace File 🗗 or Creator 🗗) with some example Wikipedia pages. Cl			
Publisher	Random House	Read more about mapping	Wikipedia templates.		dbo:publisher	dbr:Random House	
Publication date	January 8, 2013	Template Mapping (help)			dbo:releaseDate	 2013-01-08 (xsd:date) 	
Publication place	United States	map to class	Book		<u>ubo:releaseDate</u>	 2013-01-08 (xsd:date) 	
Media type	Frint (hardcover)						
Pages	208	Mappings					
ISBN	0812993802						
OCLC	785558855 🖉	Property Mapping (help)					
		template property	name				
		ontology property	foaf:name				
		Property Mapping (help)					
		template property	genre				
		ontology property	literaryGenre	/			
		Property Mapping (help)					
	(template property	pages				
l le	niversity of Mann		numberOfPages		/ersion 16.1	10 2024	22
UI	inversity of Iviau				CISIOII 10.	10.2024	~~
4							



щ°Гн

YAGO



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

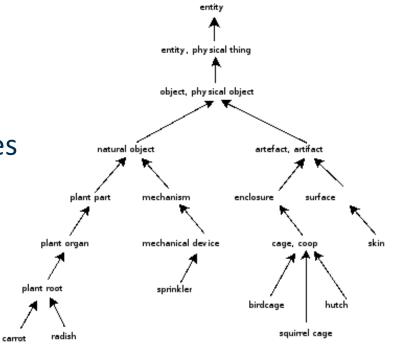
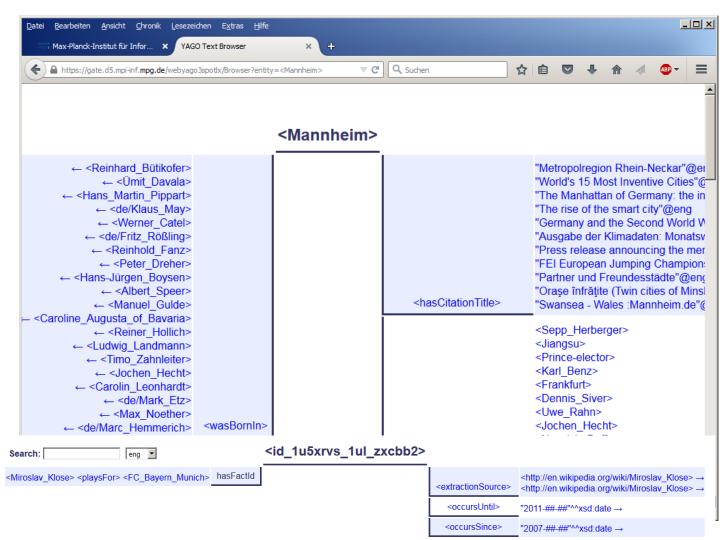


Figure 1. "is a" relation example

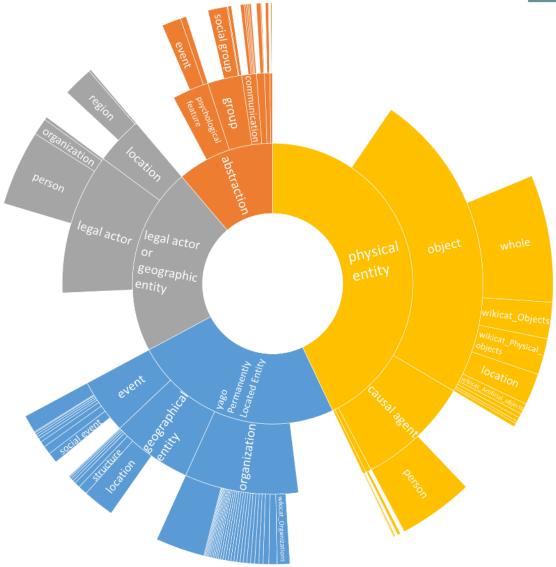
YAGO











Knowledge Graph Creation



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

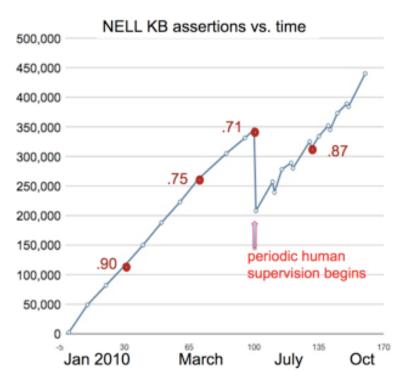


Knowledge Graph Creation: NELL



• The 2010s

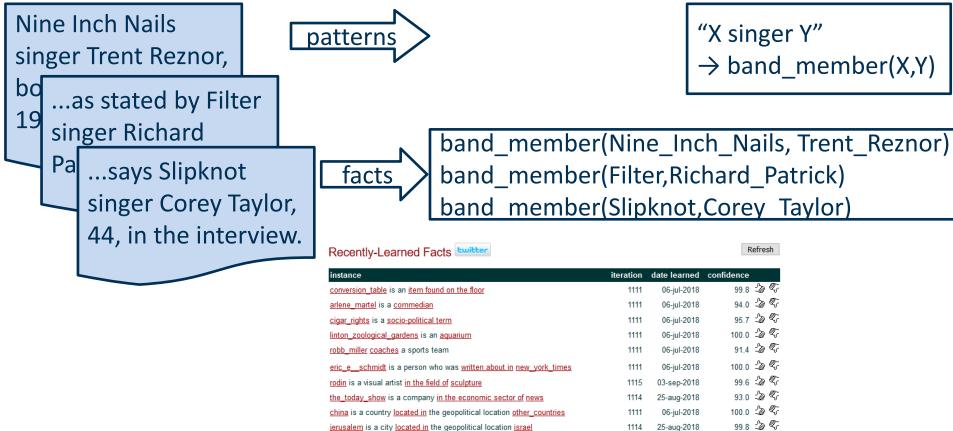
- NELL: Never ending language learner
- Input: ontology, seed examples, text corpus
- Output: facts, text patterns
- Large degree of automation, occasional human feedback
- Until 2018
 - Continuously ran for ~8 years
 - New release every few days



Knowledge Graph Creation: NELL



• Extraction of a Knowledge Graph from a Text Corpus



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

 \rightarrow 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)

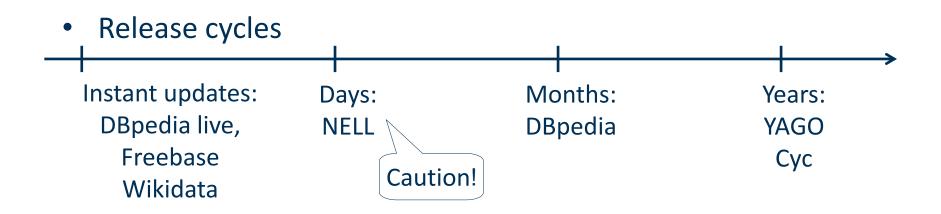


However, we usually know only little about them

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.202433See: Noy et al. (2019): Industry-scale Knowledge Graphs: Lessons and Challenges: Five diverse technology companies show how it's done

Comparison of Knowledge Graphs





• Size and density

2017

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

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Ringler & Paulheim: One Knowledge Graph to Rule them All? KI 2017

Comparison of Knowledge Graphs



2020		DBpedia	YAGO	Wikidata	BabelNet
	# Instances	5,044,223	6,349,359	52,252,549	7,735,436
	# Assertions	854,294,312	479,392,870	732,420,508	178,982,397
	Avg. linking degree	21.30	48.26	6.38	0.00
	Median ingoing edges	0	0	0	0
	Median outgoing edges	30	95	10	9
	# Classes	760	819,292	2,356,259	6,044,564
	# Relations	1355	77	6,236	22
	Avg. depth of class tree	3.51	6.61	6.43	4.11
	Avg. branching factor of class tree	4.53	8.48	36.48	71.0
	Ontology complexity	SHOFD	SHOIF	SOD	SO
		Cyc	NELL	CaLiGraph	Voldemort
	# Instances	122,441	5,120,688	7,315,918	55,861
	# Assertions	2,229,266	60,594,443	517,099,124	693,428
	Avg. linking degree	3.34	6.72	1.48	0
	Median ingoing edges	0	0	0	0
	Median outgoing edges	3	0	1	5
	# Classes	116,821	1,187	755,963	621
	# Relations	148	440	271	294
	Avg. depth of class tree	5.58	3.13	4.74	3.17
	Avg. branching factor of class tree	5.62	6.37	4.81	5.40
	Ontology complexity	SHOIFD	SROIF	SHOD	SH

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Heist, Hertling, Ringler & Paulheim: Knowledge Graphs on the Web – an Overview. 2020

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



	D	Y	W	0	Ν
Person					
Politician					
Athlete					
Actor					
Government Org.					
Company					
Political Party					
Place					
Settlement					
Country					
Work					
Album					
Song					
Movie					
Book					
Car					
Ship					
Spacecraft					
Event					
Military Conflict					
Societal Event					
Sports Event					
Chemical Substance					
Astronomical Obj.					
Planet					

	DYWO	Ν
Person		
Politician		
Athlete		
Actor		
Government Org.		
Company		
Political Party		
Place		
Settlement		
Country		
Work		
Album		
Song		
Movie		
Book		
Car		
Ship		
Spacecraft		
Event		
Military Conflict		
Societal Event		
Sports Event		
Chemical Substance		
Astronomical Obj.		
Planet		

	D	Y	W	0	N
Person					
Politician					
Athlete					
Actor					
Government Org.					
Company					
Political Party					
Place					
Settlement					
Country					
Work					
Album					
Song					
Movie					
Book					
Car					
Ship					
Spacecraft					
Event					
Military Conflict					
Societal Event					
Sports Event					
Chemical Substance					
Astronomical Obj.					
Planet					

1k

(c) Average outdegree

(a) Number of instances

3M

0

(b) Average indegree

70k

0

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Ringler & Paulheim: One Knowledge Graph to Rule them All? KI 2017 0

Comparison of Knowledge Graphs

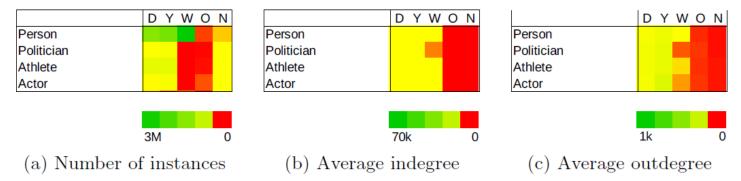


- Summary findings:
 - Persons: more in Wikidata (twice as many persons as DBpedia and YAGO)
 - Countries: more details in Wikidata
 - Places: most in DBpedia
 - Organizations: most in YAGO
 - Events: most in YAGO
 - Artistic works:
 - Wikidata contains more movies and albums
 - YAGO contains more songs

Caveats



• Reading the diagrams right...

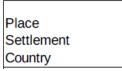


- So, Wikidata contains more persons
 - but less instances of all the interesting subclasses?
- There are classes like Actor in Wikidata
 - but they are hardly used
 - rather: modeled using *profession* relation

Caveats

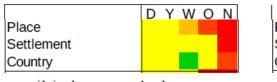


Reading the diagrams right... (ctd.)





(a)) Number	of instances	
-----	----------	--------------	--







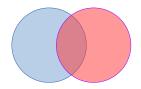
(c) Average outdegree

- 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

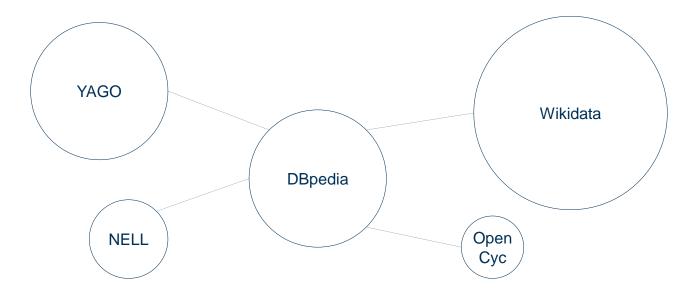
University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024



- 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



University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Ringler & Paulheim: One Knowledge Graph to Rule them All? KI 2017



- 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





42



- 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



• 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 $|C| = |F| * P * \frac{1}{R}$

• The existing links are used to compute precision and recall

Overlap of Knowledge Graphs

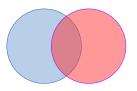
Mathematical derivation:

- 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 $|C|$

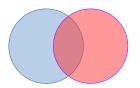
$$|C| = |F| * P * \frac{1}{R}$$







- Experiment:
 - We use the same 25 classes as before

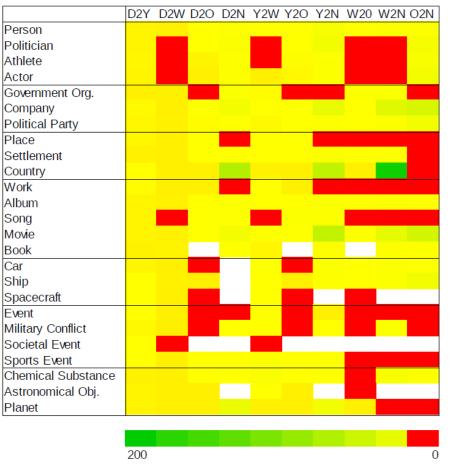


- Measure 1: overlap relative to smaller KG (i.e., potential gain)
- Measure 2: overlap relative to explicit links (i.e., importance of improving links)
- Link generation with 16 different metrics and thresholds
 - Intra-class correlation coefficient for |C|: 0.969
 - Intra-class correlation coefficient 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
 University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024
 Ringler & Paulheim: One Knowledge Graph to Rule them All? KI 2017



	D2Y	D2W	D20	D2N	Y2W	Y2O	Y2N	W20	W2N	O2N
Person										
Politician										
Athlete										
Actor										
Government Org.										
Company										
Political Party										
Place										
Settlement										
Country										
Work										
Album										
Song										
Movie										
Book										
Car										
Ship										
Spacecraft										
Event										
Military Conflict										
Societal Event										
Sports Event										
Chemical Substance										
Astronomical Obj.										
Planet										



(a) Overlap as potential gain

1

(b) Overlap relative to existing links

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Ringler & Paulheim: One Knowledge Graph to Rule them All? KI 2017

0



- 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



• There are quite a few metrics for evaluating KGs

- size, degree, interlinking, quality, licensing, ...

Dimension	Abr	Metric	Description	Туре
	A1	accessibility of the SPARQL end- point and the server	checking whether the server responds to a SPARQL query [18]	QN
Availability	A2	accessibility of the RDF dumps	checking whether an RDF dump is provided and can be down- loaded [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 re- turned (iii) that useful data (particularly RDF) is returned upon lookup of a URI, (iii) for changes in the URI i.e the compli- ance 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 de- scription 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	11	detection of good quality inter- links	(i) detection of (a) interlinking degree, (b) clustering coefficient, (c) centrality, (d) open sameAs chains and (e) description richness through sameAs by using network measures [25], (ii) via crowdsourcing [1,65]	QN
	12	existence of links to external data providers	detection of the existence and usage of external URIs (e.g. us- ing owl:sameAs links) [31]	QN
	13	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	S 1	usage of digital signatures	by signing a document containing an RDF serialization, a SPARQL result set or signing an RDF graph [13,18]	QN
	S 2	authenticity of the dataset	verifying authenticity of the dataset based on a provenance vo- cabulary such as author and his contributors, the publisher of the data and its sources (if present in the dataset) [18]	QL
D.C.	P1	usage of slash-URIs	checking for usage of slash-URIs where large amounts of data is provided [18]	QN
Performance	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 re- quests divided by ten is not longer than the time it takes to an- swer one request [18]	QN

Table 2

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	mFreq	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
	mextVoc	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
	mAvai	0.9961	0.9998	1	0.9999	0.7306	2
	<i>m_{SPARQL}</i>	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	$m_{macLicense}$	1	0	0	1	0	1
Interlinking	mInst	0.592	0.018	0.443	0	0.305	2
	muRIs	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	

Table 14

work with an example weighting which would be reasonab

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

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

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024



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



Some back of the envelope calculations...

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Paulheim: How much is a triple? Estimating the Cost of Knowledge Graph Creation, 2018

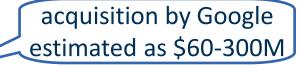
- Case 1: manual curation
 - Cyc: created by experts
 Total development cost: \$120M
 Total #statements: 21M
 - ightarrow \$5.71 per statement
 - Freebase: created by laymen
 Assumption: adding a statement to Freebase
 equals adding a sentence to Wikipedia
 - English Wikipedia up to April 2011: 41M working 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
 - ightarrow \$2.25 per statement

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

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Paulheim: How much is a triple? Estimating the Cost of Knowledge Graph Creation, 2018









 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

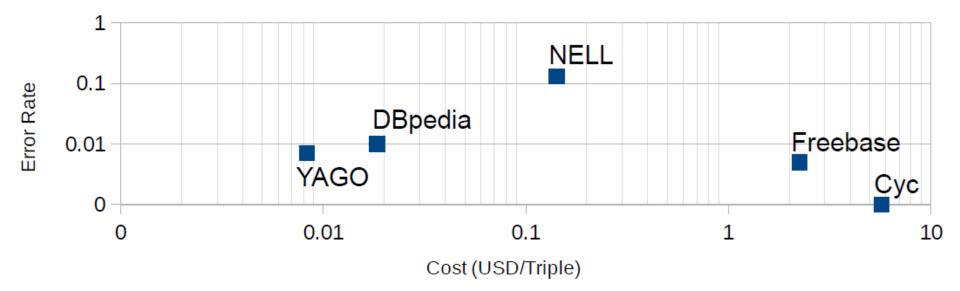








- we can pay for accuracy
- NELL is a bit of an outlier





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.

- <u>README.txt</u>
- <u>hetrec2011-movielens-2k.zip</u>

https://grouplens.org/datasets/



Enhancing the Coverage of Knowledge Graphs

Richard Devine

• B12

Cabaret Voltaire



Delicious Bookmarks Only existing pages have 105,000 bookmarks from 1867 users categories README.txt hetrec2011-delicious-2k.zip Lists may also link to non-existing pages Last.FM List of intelligent dance music artists 92,800 artist listening records from 1892 users. From Wikipedia, the free encyclopedia README.txt This section does not cite any sources. Please help improve this section by adding of hetrec2011-lastfm-2k.zip Find sources: "List of intelligent dance music artists" - news · newspapers · books · scholar · JSTOR This is a list of notable music artists who play intelligent dance music (IDM) genre MovieLens + IMDb/Rotten Tomatoes Contents [hide] 1 # 2 A-K 86,000 ratings from 2113 users. 3 L-Z 4 References README.txt hetrec2011-movielens-2k.z # [edit] 808 State µ-ziq https://grouplens.org/datasets/ A-K [edit] Actress Benn Jordan Casino Versus Japan Dopplereffekt Funkstörung Jan Jelinek Acustic Biosphere Ceephax Acid Crew Chris Douglas The Future Sound Of I Jega ndon Air Liquide Björk^[1] Cex Drexciva Gas Jello Christ Eight Frozen Modules Alarm Will Sound The Black Dog Gescom Jlin John Tejada Alva Noto Blanck Mass Chris Clark Emptyset Global Communication Amon Tobin · Boards of Canada Clocolan Esem Global Goon Jon Hopkins^[3] Goldie Andy Stott Bochum Welt Cylob FaltvDL Kettel Aphex Twin Boom Bip Daedelus Fennesz Zachary Grav^[2] Kevin Blechdom Apparat Brothomstates Deadbeat The Field Gridlock Kid606 Arovane Burial Deepchord The Flashbulb Himuro Yoshiteru Kodomo Demdike Stare Koreless^[4] Atypic Bvdub Floating Points Kim Hiorthøy Flying Lotus Autechre C418 Deru . I am Robot and Proud

Eorest Swords

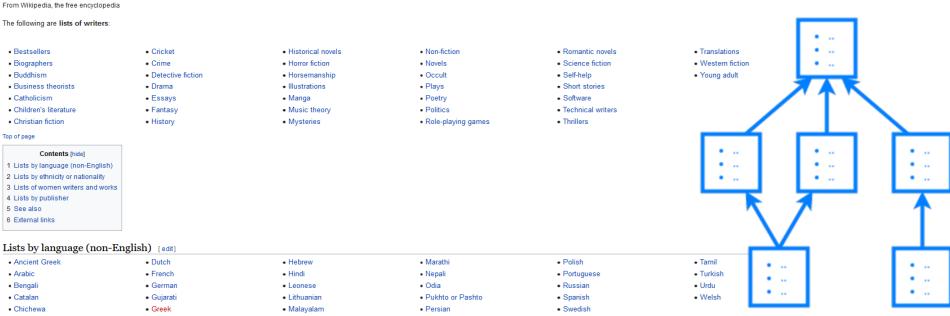
Innovaders

Entity Extraction from List Pages



• Lists form (shallow) hierarchies

Lists of writers



Top of page

Lists by ethnicity or nationality [edit]

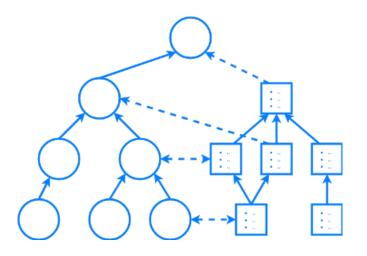
 African writers 	 Barbadian writers 	 English writers 	 Irish writers 	 Nepali writers 	 Romanian writers
 African-American writers 	 Beninese writers 	 Egyptian writers 	 Italian writers 	 New Zealand writers 	 Russian authors
 Albanian writers 	 Black British writers 	 Georgian writers 	 Jewish American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

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024

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

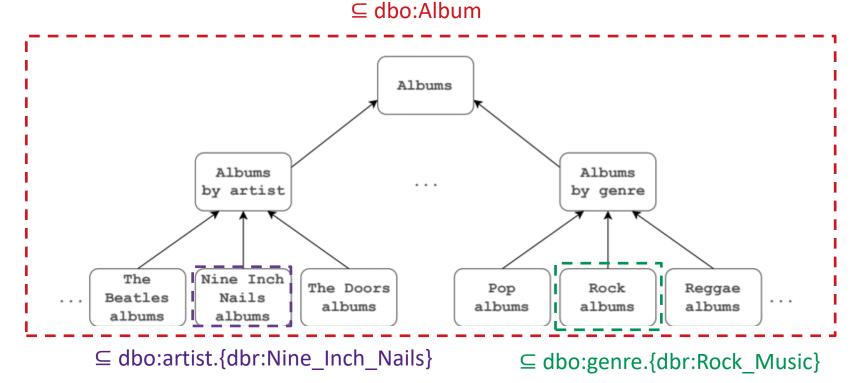
- Patricia Aakhus (1952–2012), The Voyage of Mael Duin's Curragh
- Atia Abawi
- Edward Abbey (1927–1989), The Monkey Wrench Gang
- Lynn Abbey (born 1948), Daughter of the Bright Moon
- Belle Kendrick Abbott (1842-1893), Leah Mordecai
- Eleanor Hallowell Abbott (1872–1958), poet, novelist and short story 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



University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Heist & Paulheim (2019): Uncovering the Semantics of Wikipedia Categories

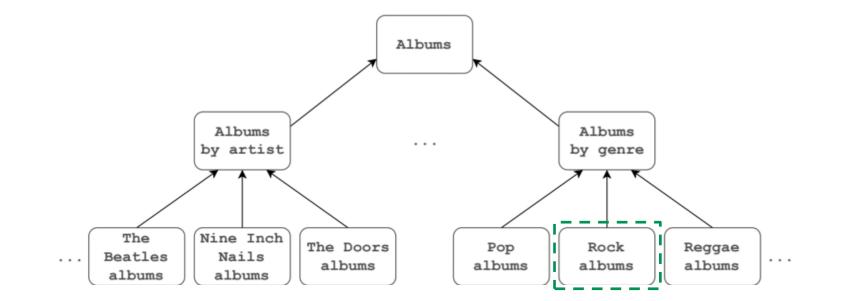
щля

UNIVE

OF MA

Data and Web Science Group

Cat2Ax: Axiomatizing Wikipedia Categories



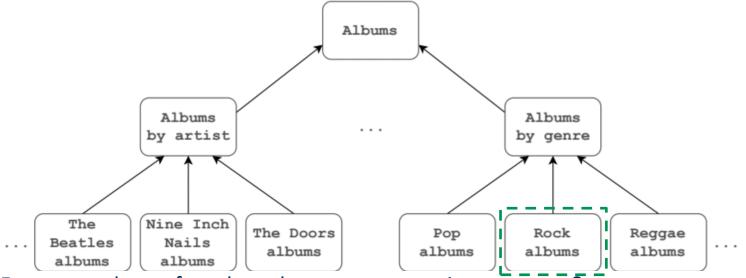
⊆ dbo:genre.{dbr:Rock_Music} ⊆ dbo:genre.{dbr:Rock_(Rapper)}

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Heist & Paulheim (2019): Uncovering the Semantics of Wikipedia Categories



Cat2Ax: Axiomatizing Wikipedia Categories



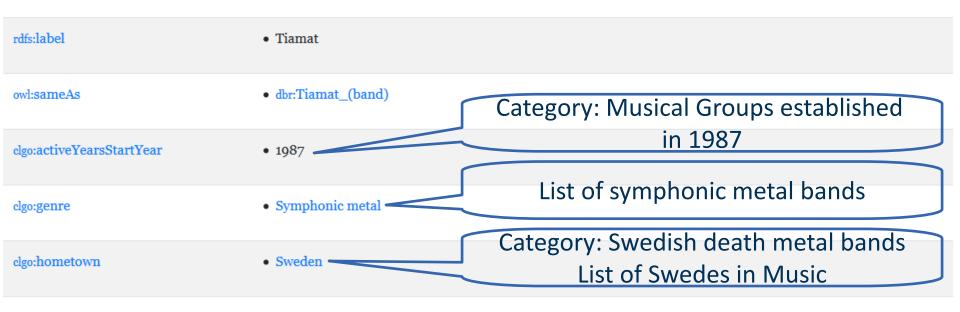


- Frequency: how often does the pattern occur in a category?
 - i.e.: share of instances that have dbo:genre.{dbr.Rock_Music}?
- Lexical score: likelihood of term as a surface form of object
 - i.e.: how often is Rock used to refer to dbr:Rock_Music?
- Sibling score: how likely are sibling categories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

i.e., are there sibling categories with a high score for dbo:genre?
 University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024
 Heist & Paulheim (2019): Uncovering the Semantics of Wikipedia Categories

CaLiGraph Example





Pushing Entity Coverage Further



• Beyond red links (2020)

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
- Volker Borchert drums (1992, 2015-2016)

Discography [edit]

Albums [edit]

- Stahlwerksynfonie (1981)
- Volle Kraft Voraus! (1982)
- Entering the Arena (1985)
- / (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)

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024

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 Luliaby 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) ∃topSection.{"Discography" ⊑ ∃artist.{>PageEntity<}/

(type, MusicalWork) ∃topSection.{"Discography"} ⊑ MusicalWork

(musicalBand, Guns_N'_Roses)

∃topSection.{"Discography"} □ sectionEntityType.{Band} ⊑ ∃musicalBand.{>SectionEntity<}

Page Entity Axl Rose From Wikipedia, the free encyclopedia Discography Jop Section Section with Guns N' Roses 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 Stradin) 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 (edt) 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

Voice

Voice

New Looney Tunes (TV show)[123]

Jimmy Kimmel Live!

Scooby-Doo and Guess Who? (Tv Show)

2012

2018

2021

Himself

Himself

Himself

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



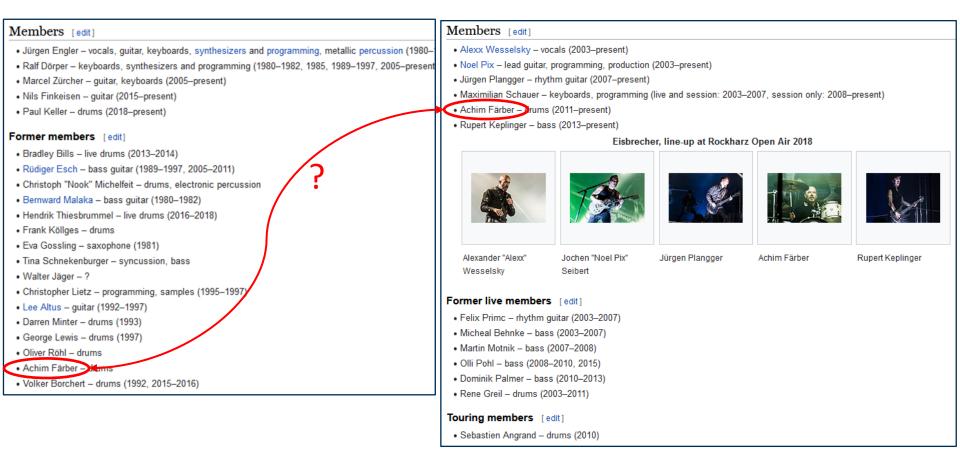
- Latest version 2.1
 - 15M entities
 - incl. 8M from listings
 - Caveat:
 - Disambiguation!



Entity Disambiguation



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



CaLiGraph Glitches



CaLiGraph

🖿 Formats 🥆

🕑 Sparql Endpoint

About: clgr:Mannheim 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 local Value Property 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 places included are towns, the list also comprises villages, cities, districts, counties, etc. with similar links. rdfs:label Mannheim 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] Chişinău^{[3][4]} Moldova clgo:country • Alba Iulia, Romania Kiev, Ukraine • 🕂 Tbilisi, Georgia Germany C Ankara, Turkey Mannheim, Germany • 💽 Tel Aviv, Israel • Bucharest, Romania Minsk, Belarus • 🗾 Vilnius, Lithuania • E Chernivtsi, Ukraine Odessa, Ukraine • Merevan, Armenia • Grenoble, France • Reggio Emilia, Italy Planned capi rdf:type • Iași, Romania Sacramento, United States City in Baden Comrat Twin town or Bălți, Moldova^[1] Coat of arms Criuleni • Jurbarkas, Lithuania^[5] French exon • Twin town or sister city in Lithuania University town in Germany owl:NamedIndividual • City or town in Germany • Most polluted city in the world

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





Ar Recent Wiki Activity

Amar'e Stoudemire

Carmelo Anthony

Frank Ntilikina • 1 minute ago

)g in

Q

170p

One (but not the only!) possible source of coverage bias \bullet

> POPULAR PAGES - COMMUNITY - COMPULAR PAGES - COMPU in: HasTwitter HasMvSpace, HasFacebook, and 2 more

The purpose of this wiki is the same that of Wikipedia, it is to create an encyclopedia which

is a comprehensive summary of information

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

Did vou ever have a Wikipedia article vou

comprehensive knowledge, because Wikipedia

worked on deleted and did not have a backun?

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

deleted on Wikipedia. Like that you can still go into the history of an article and find your edits.

Did you feel frustrated like me? You are in the

together with Wikipedia we will have

deletes so much

right place

Main Page

rescued 53 934 articles

Some bot code is available at GitHub d.

How does this work?

You should be able to actually sign up and edit articles

articles are often relisted for deletion again soon

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

from all branches of knowledge. The only difference between the Wikipedia and

Speedy Deletion Wikia Main Page

Articles about long-tail entities become deleted

Page Discu

Wikipedia:Deletion policy

From Wikipedia, the free encyclopedia

1 Reasons for deletion

2 Alternatives to deletion 2.1 Editing and discussion

2.2 Tagging

2.3 Merging

2.4 Redirection

2.5 Incubation 2.6 Other projects

2.7 Archiving

3.1 Copyright violations

3.2 Speedy deletion

3.3 Proposed deletion

3.4 Deletion discussion

3.5 Page deletion

3 Processes

"WP:DP", "WP:DELETE", and "WP:DEL" redirect here. For other uses, see Wikipedia:Deletion pro-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 previou reverse ("undelete") any deletion. All such actions (other than viewing) ar administrators normally will not delete it

Contents [hide]

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

What links here



3.3.1 Proposed deletion of biographies of living people

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,

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

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

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

http://speedydeletion.wikia.com



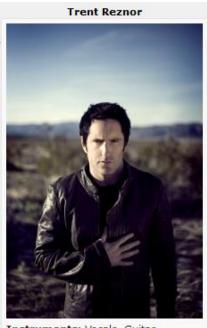


- Why stop at Wikipedia?
- Wikipedia is based on the MediaWiki software
 - …and so are thousands of Wikis
 - Fandom by Wikia: >385,000 Wikis on special topics
 - WikiApiary: reports >20,000 installations of MediaWiki on the Web





• Collecting Data from a Multitude of Wikis



Instruments: Vocals, Guitar, Keyboards, Bass, Marimba, Saxophone, Small Percussion Years: 1988-present Tours: VIVIsectVI-present



Role Composer Born May 17, 1965 Mercer, Pennsylvania, USA



Born May 17, 1965 New Castle, Pennsylvania, United States

Other David Lynch Projects Lost Highway (Soundtrack - "Videodrones; Questions," "Driver Down") "Came Back Haunted" (Music video)



- 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?





• 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 Griffindors 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

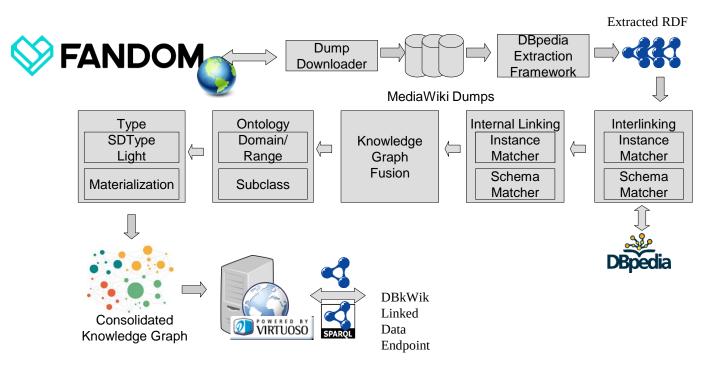
University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 <u>http://dbkwik.org/</u>



- 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



- Heuristics
 - Ontology induction
 - Instance/Schema Matching



University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 Hertling & Paulheim: DBkWik: A Consolidated Knowledge Graph from Thousands of Wikis. ICBK 2018



- 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)



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

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.202482Hertling & Paulheim: WebIsALOD: Providing Hypernymy Relations extracted from the Web as Linked Open Data. ISWC 2017

WeblsALOD

- 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 China	Beijing	20,693,000 ^[1]	2012	1.529
2	💶 India	New Delhi	16,787,949[2]	2014	0.90%
3	Japan	Токуо	13,189,000 ^[3]	2011	10.329
4	Market Philippines	Manila	12,877,253 ^[4]	2015	12.449
5	Russia	Moscow	11,541,000 ^[5]	2011	8.079
6	Egypt	Cairo	10,230,350	2012	11.109
7	- Indonesia	Jakarta	10,187,595 ^[6]	2011	4.189
8	Democratic Republic of the Congo	Kinshasa	10,125,000[7]	2012	12.309
9	South Korea	Seoul	9,989,795 ^[8]	2015	20.479
10	Bangladesh	Dhaka	8,906,000 [9]	2011	5.561
11	Mexico	Mexico City	8,851,080 ^[10]	2010	7.519
12	💳 Iran	Tehran	8,846,782	2014	9.91
13	Inited 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	C 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	C Spain	Madrid	3,233,527[22]	2012	6.84
28	North Korea	Pyongyang	3,144,005		12.63
29	S Afghanistan	Kabul	3,140,853		10.28
30	Kenya	Nairobi	3,138,369	2010	7.67



WeblsALOD



- Extraction of type information using Hearstlike 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

Common Crawl

400M hypernymy relations

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 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

Narrower concepts

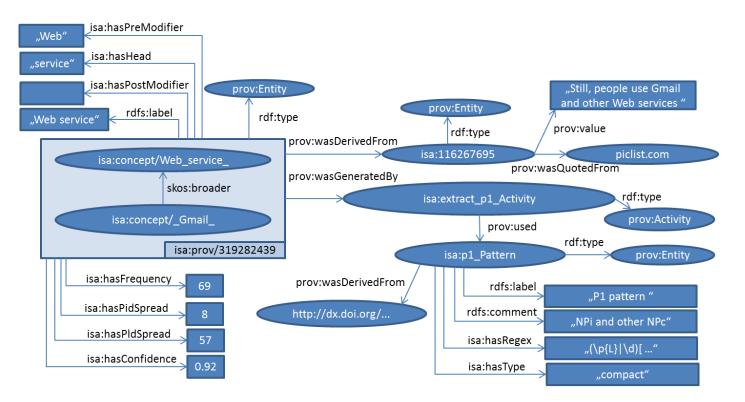
label	provenance	confidence	label	label provenance
	provenance	conndence	laber	label provenance
writer	isap:391280092	0.799331	george orwell	george orwell isap:386468501
great idea	isap:493244047	0.672180	science fiction	science fiction isap:275868279
several magazine	isap:104101164	0.655684	franz kafka	franz kafka isap:159147340
category	isap:104762336	0.477191	steve almond	steve almond isap:392552636
artist	isap:387107910	0.471280	dan brown	dan brown isap:157209267
blog	isap:492616562	0.458511	james joyce	james joyce isap:159394667
writers	isap:439522913	0.427701	stephen king	stephen king isap:306753456
story	isap:122402598	0.306667	flannery oconnor	flannery oconnor isap:266705231
group	isap:115379219	0.299656	alice munro	alice munro isap:162537618
poet	isap:492284397	0.287519	ayn rand	ayn rand isap:301402665

University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.2024 <u>http://webisa.webdatacommons.org/</u>

WebIsALOD



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

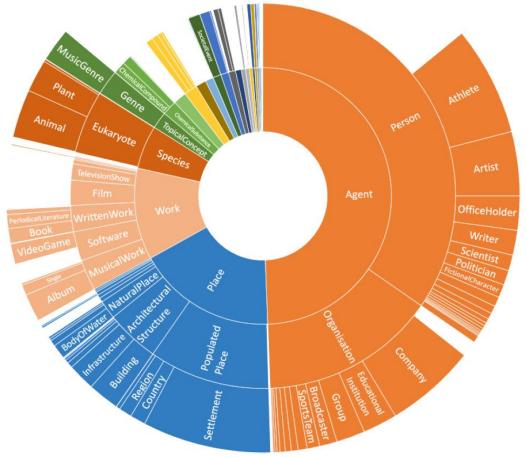


University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.202485Hertling & Paulheim: WebIsALOD: Providing Hypernymy Relations extracted from the Web as Linked Open Data. ISWC 2017

WebIsALOD



• Estimated contents breakdown



University of Mannheim | IE650 Knowledge Graphs | Public Knowledge Graphs | Version 16.10.202486Hertling & Paulheim: WebIsALOD: Providing Hypernymy Relations extracted from the Web as Linked Open Data. ISWC 2017

WeblsALOD



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

WeblsALOD



- 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)

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?



