Complex Question Answering

over Knowledge Graphs

Mikhail Galkin
PhD, Research Scientist
TU Dresden & Fraunhofer IAIS
Dresden

@migalkin @michael_galkin @mgalkin

Outline

- Introduction: Knowledge Graphs
- Large KGs:
  - Template-based
  - Pipeline-based
- Smaller KGs: Neural Reasoning
  - Neural and Multi-Hop QA
  - Query Embedding

Fraunhofer
Dresden IAIS

About

SPEAKER

Industrial Speech Assistance Platform

- ➔ ConvAI
- ➔ Question Answering
- ➔ Knowledge Graphs
- ➔ Speech
- ➔ Privacy
- ➔ Customizable & Domain-specific

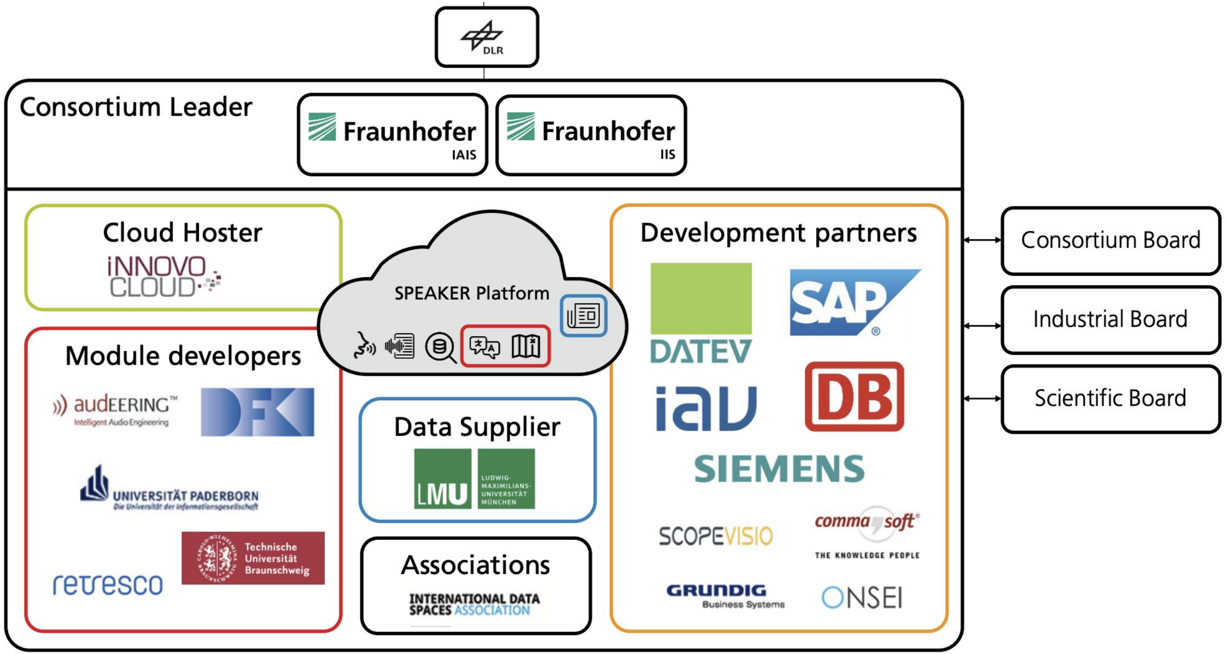

On the definition of a Knowledge Graph

Given entities E, relations R, KG is a directed multi-relational graph G that comprises triples (s, p, o)

$$\mathcal{G} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$$

$$(s, p, o) \in \mathcal{G}$$

* describes entities and relations
* defines a schema
* interrelating arbitrary entities
* various topical domains

"Every RDF / LPG / RDF* graph is a knowledge graph"

"Abstract schema and instances"

Graph-structured world model

World models?

Entities and relations define our **domain of discourse**

How to encode it?

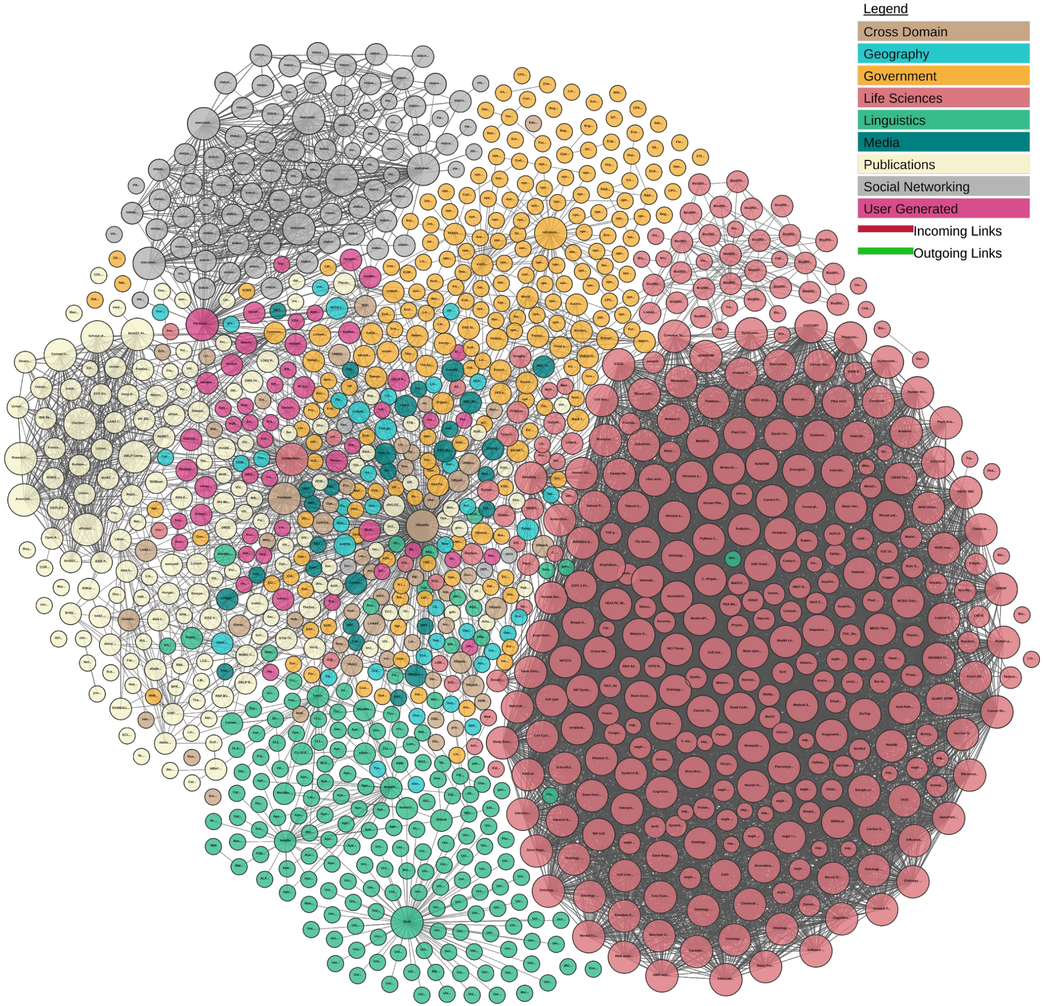

Source: https://lod-cloud.net/

On representation of Knowledge Graphs

Symbolic

```
s,p,o
p(s,o)
(s,p,o)
```

Vector

$$s, p, o \in \mathbb{R}^d$$

On representation of Knowledge Graphs

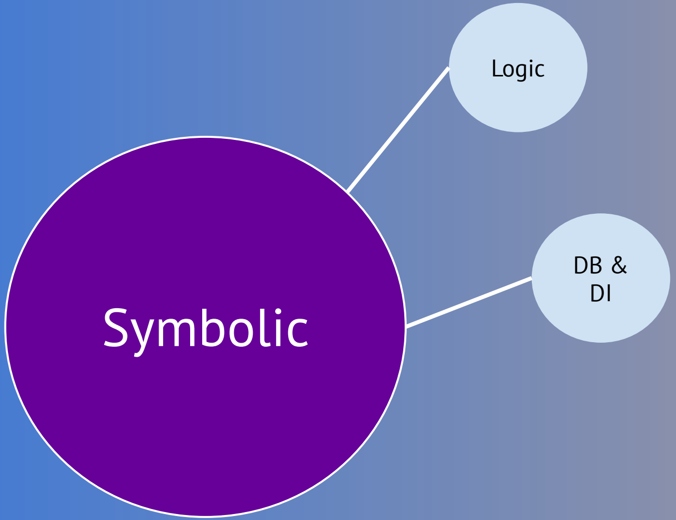

Open-world assumption

Closed-world assumption

Temporal / evolving

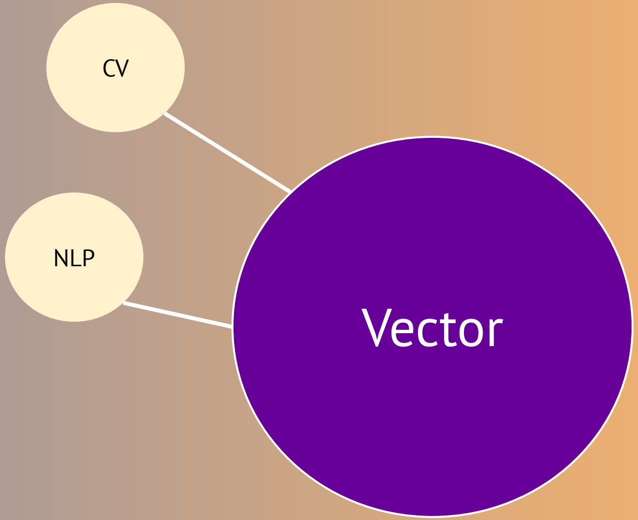

s,p,o
p(s,o)
(s,p,o)

$$s, p, o \in \mathbb{R}^d$$

Symbolic: Triples

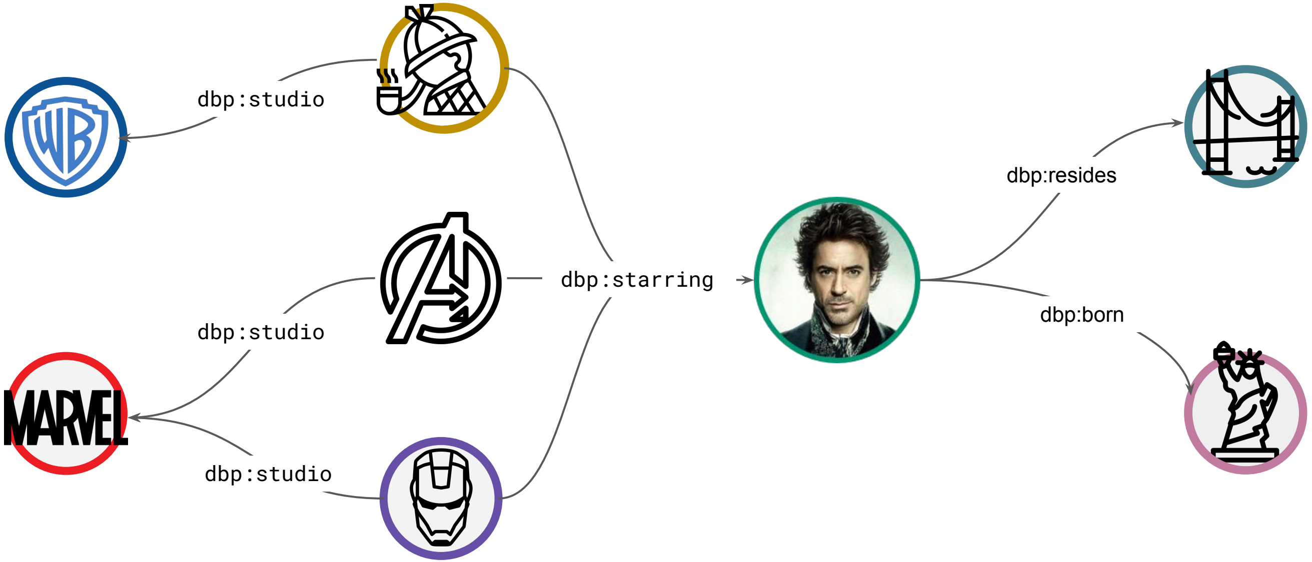

RDJ	dbp:resides	SF .	Avengers	dbp:studio	Marvel .
RDJ	dbp:born	NY .	Avengers	dbp:starring	RDJ .
Sherlock_Holmes	dbp:studio	WB .	Iron_Man	dbp:studio	Marvel .
Sherlock_Holmes	dbp:starring	RDJ.	Iron_Man	dbp:starring	RDJ .

Symbolic: Description Logics

Based on logical formalisms, e.g., Description Logics (DL), RDFS, OWL

schema, ontology, theory	**TBox**	$\mathtt{SuccessfulArtist} \sqsubseteq \geq 1\ \mathtt{actedIn.Blockbuster}$
instances, facts, assertions	**ABox**	$\mathtt{SuccessfulArtist(RobertDowneyJr)}$
restrictions, constraints	**RBox**	$\mathtt{actedIn} \sqsubseteq \mathtt{participated}$

Logically consistent collection of axioms

Vector: Embeddings

$$E \in \mathbb{R}^{N_e \times d}$$
$$R \in \mathbb{R}^{N_r \times d}$$

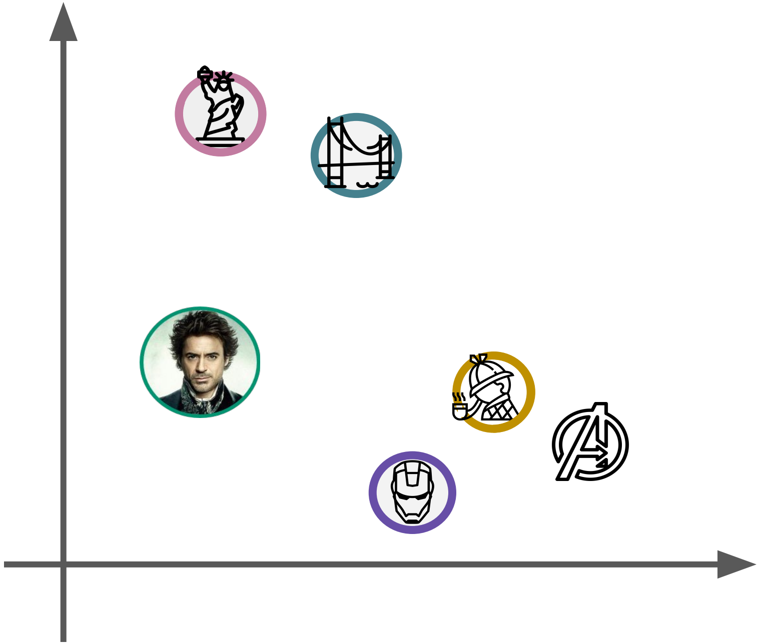

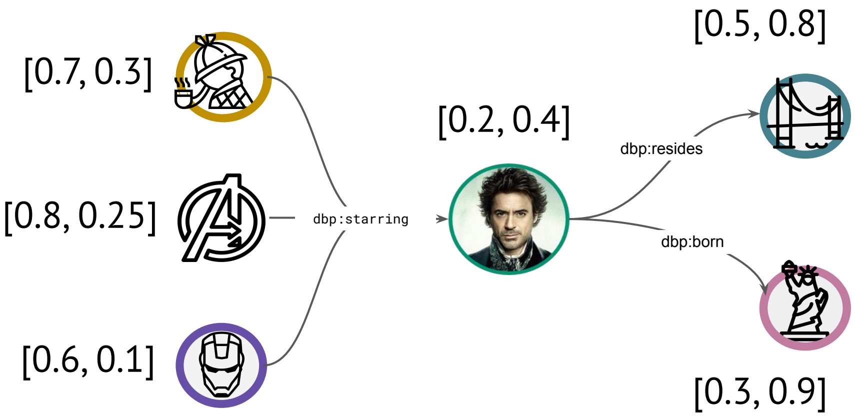

KG Embeddings: PyKEEN 1.0

PyKEEN

build passing | License MIT | DOI 10.5281/zenodo.3982977 | Optuna integrated

https://github.com/pykeen/pykeen

→ **PyTorch** 😍
→ 13 datasets + your own graphs
→ 23 KG embedding models and counting

→ 7 losses
→ 6 optimizers
→ 6 metrics
→ 5 regularizers
→ 2 training loops
→ 2 negative samplers
→ Tracking in MLFlow, WANDB

📈 Benchmarked!

Ali et al. Bringing Light Into the Dark: A Large-scale Evaluation of Knowledge Graph Embedding Models Under a Unified Framework. arxiv:2006.13365

Building KGs

Knowledge Graph

Semantic Data Integration

Structured Sources

XML, CSV, JSON

APIs

NoSQL

RDF Storage

GraphDBs

JSON

XML

SQL

CSV

Knowledge Graph

Information Retrieval & NLP

Unstructured Sources

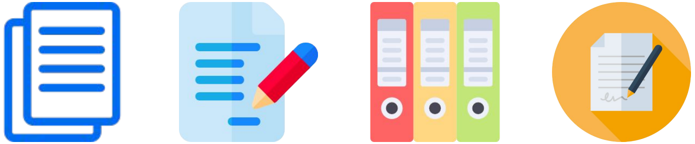

Graph Databases

RDF

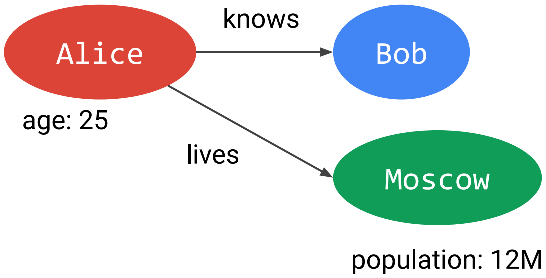

- Query language: SPARQL
- Predicate attributes only from RDFS/OWL
- Semantic schema
- Logical reasoning

LPG (Labeled Property Graph)

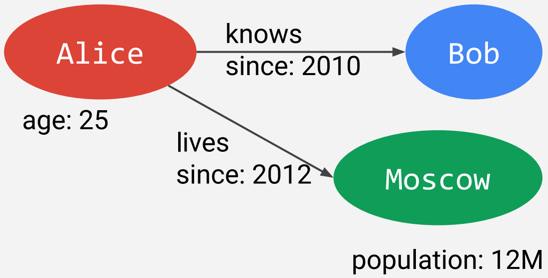

- Query languages: Cypher, Gremlin, GraphQL
- Key-value predicate attributes
- Non-semantic schema
- No reasoning

Graph Databases - Queries

SPARQL	Cypher

`SELECT ?s ?friend WHERE {`
`    ?s a :Person;`
`        :name "John" ;`
`        :knows ?friend .}`
	`MATCH (s:Person)-[:knows]-(friend)`
	`WHERE s.name = "John"`
	`RETURN s, friend ;`

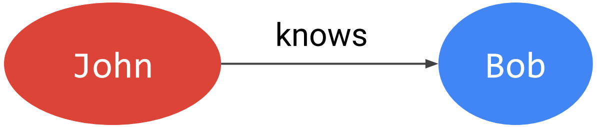

Graph Databases - Queries

SPARQL* (Reification)	Cypher
SELECT ?s WHERE { <<?s :knows :js>> :since 2001 }	MATCH (s:Person)-[:knows {since:2001}] -> (js) RETURN s;

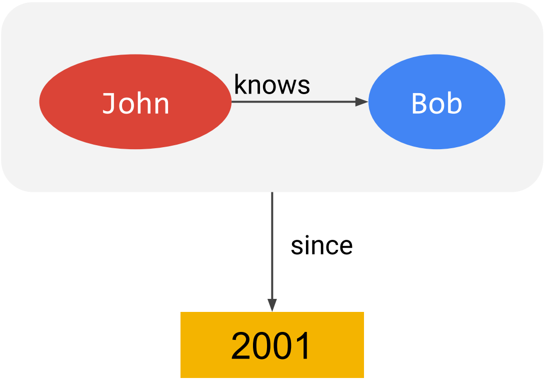

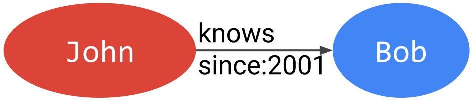

Knowledge Graphs in the Wild

https://wikidata.metaphacts.com/

Knowledge Graphs in the Wild

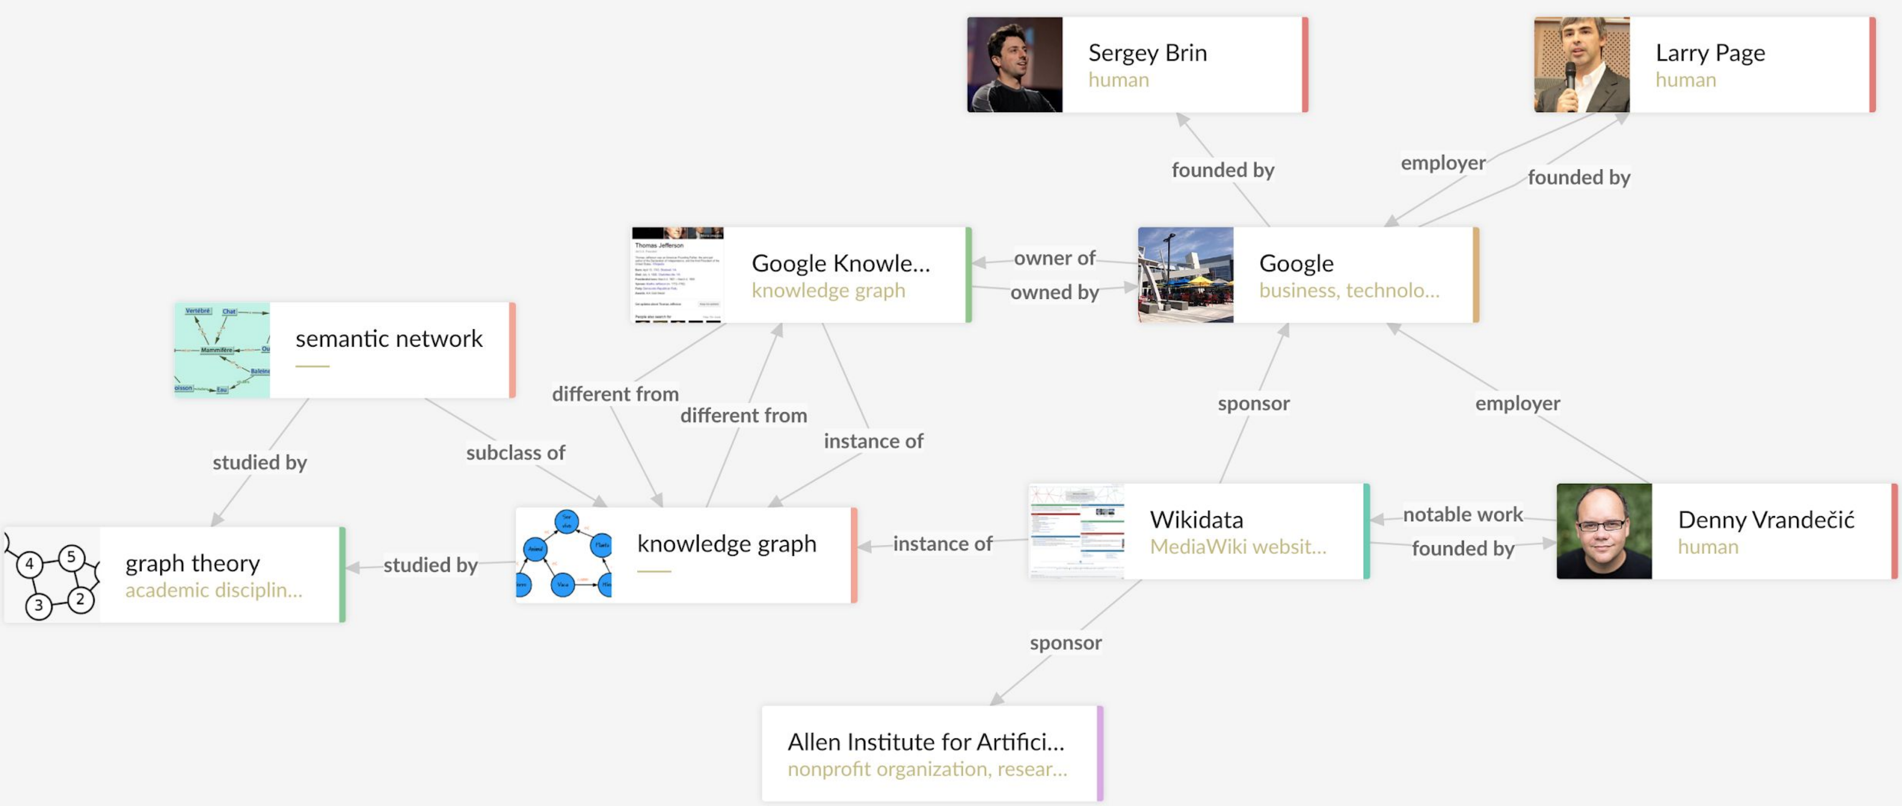

https://wikidata.metaphacts.com/

www.wikidata.org

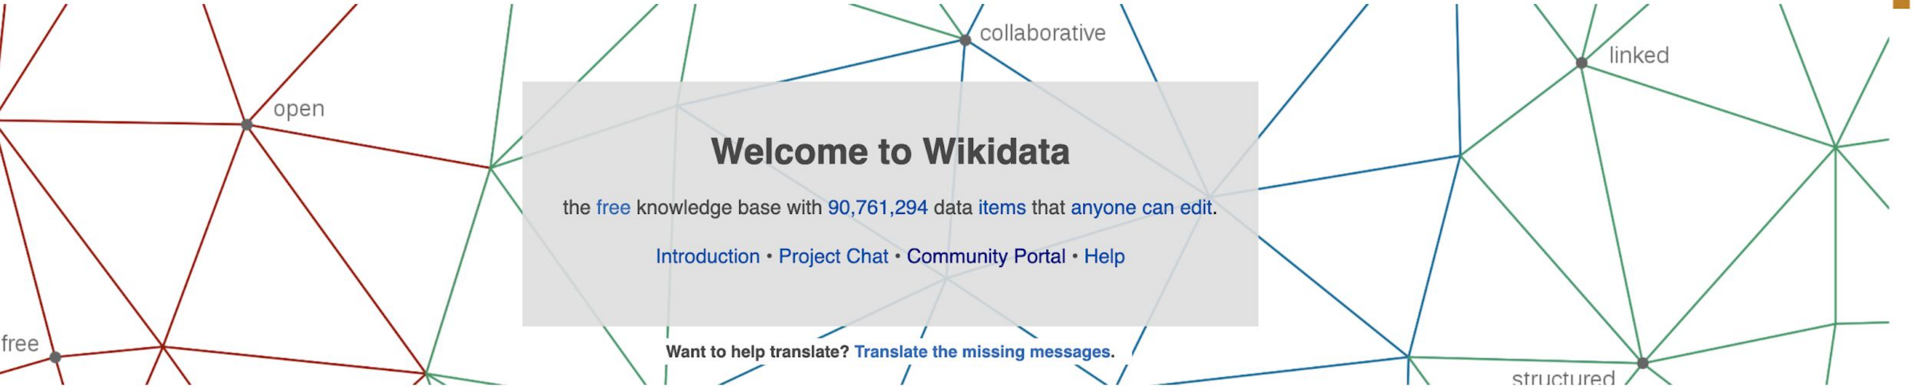

90M+	entities	**(nodes)**
1150M+	statements	**(edges)**
6K+	properties	**(edge types)**

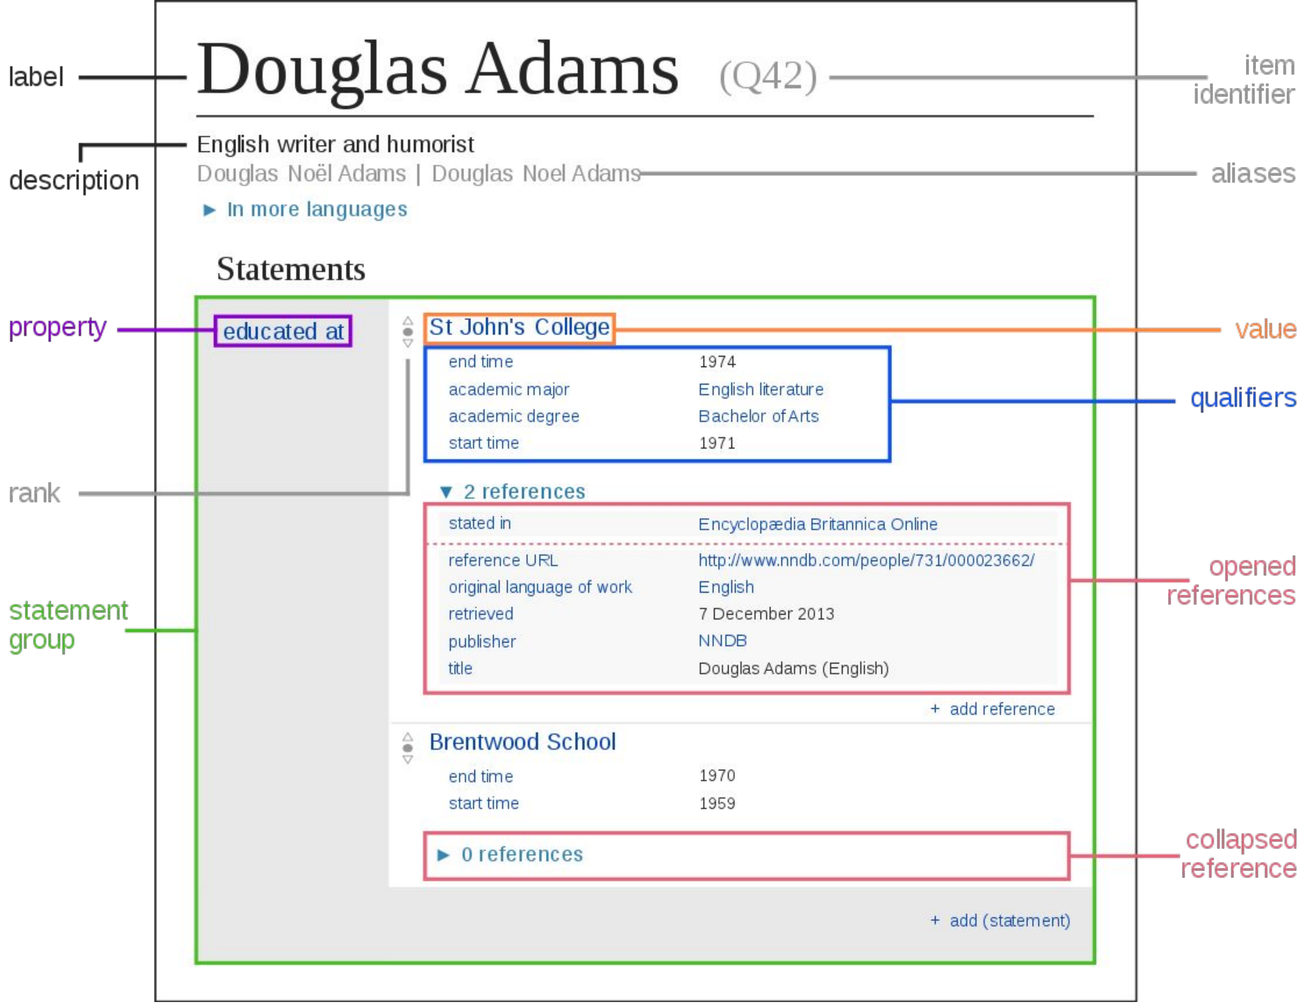

Wikidata Data Model

https://www.mediawiki.org/wiki/Wikibase/DataModel/Primer

Question Answering

MRC

Machine Reading
Comprehension

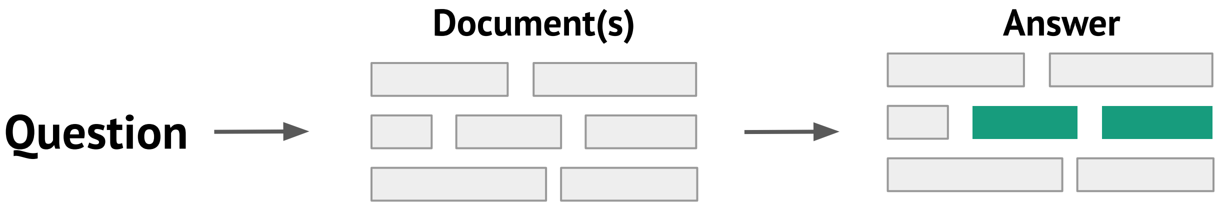

Question Answering

MRC

Machine Reading
Comprehension

Question

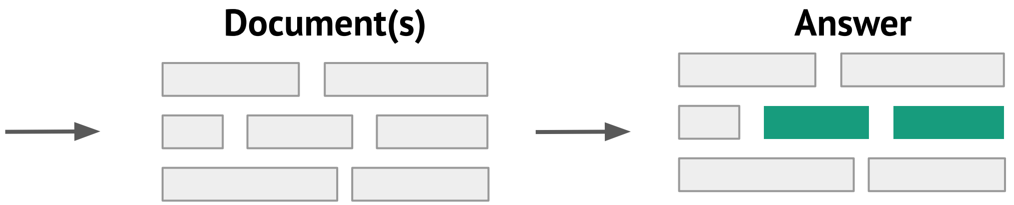

KGQA

Knowledge Graph-based
Question Answering

Question

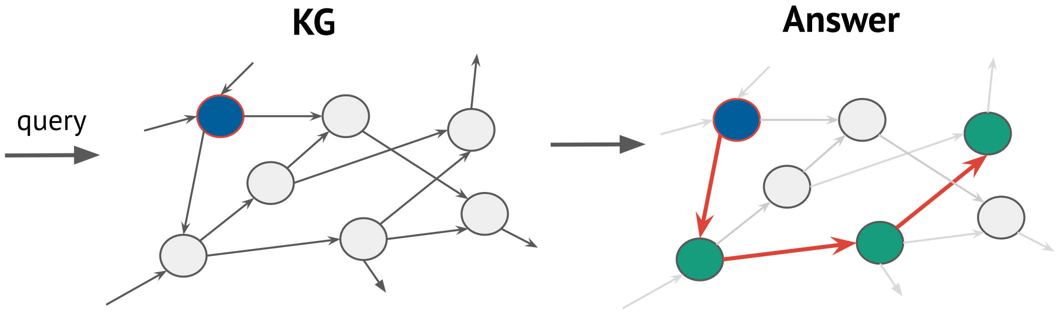

✓ Simple KGQA

Pretty much solved!

1-hop

Factoid questions

Albert Einstein

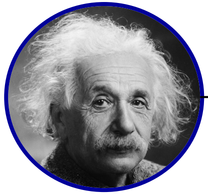

educated at

P69

?

Q937

Where was Albert Einstein educated?

Lukovnikov et al. Pretrained Transformers for Simple Question Answering over Knowledge Graphs. ISWC 2019

✓ Simple KGQA

Pretty much solved!

1-hop

Factoid questions

Albert Einstein

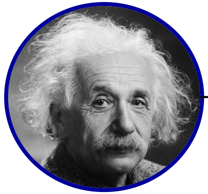

educated at

P69

?

Q937

Where was Albert Einstein educated?

University of Zurich

ETH Zurich

	Entity Span Accuracy	Avg. F1	Relation Accuracy
BiLSTM	93.2	96.7	82.4
BERT	95.2	97.5	83.5

Lukovnikov et al. Pretrained Transformers for Simple Question Answering over Knowledge Graphs. ISWC 2019

Complex KGQA

Where was the author of the theory of relativity educated?

Multi-hop

Complex

FOL / EPFO queries

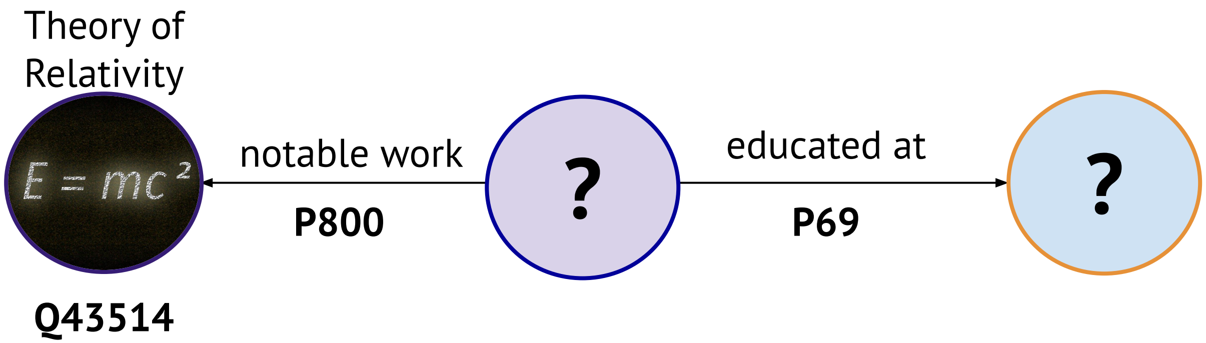

$$Q = y,\ \exists x:\ \text{notableWork}(x, \text{Theory of Relativity}) \land \text{educatedAt}(x, y)$$

Lukovnikov et al. Pretrained Transformers for Simple Question Answering over Knowledge Graphs. ISWC 2019

Complex Question Answering

Large KGs (Wikidata-scale)	Smaller KGs
>1M triples	< 1M triples
Build SPARQL queries Execute against a graph database	In-memory neural reasoning Graph embeddings
Supervised ML methods as certain pipeline components	End-to-end or mostly neural

Complex Question Answering

Large KGs (Wikidata-scale)	Smaller KGs
Pre-defined SPARQL templates	Neural Multi-Hop Reasoning
NL -> SPARQL pipelines	Query Embedding

Complex Question Answering

Large KGs (Wikidata-scale)	Smaller KGs
Pre-defined SPARQL templates	Neural Multi-Hop Reasoning
NL -> SPARQL pipelines	Query Embedding

Pre-defined SPARQL templates

Template-based QA

Natural language

`How many {Entity} did {Entity} star in?`

Pre-defined template

```
SELECT COUNT(?x) WHERE {
    ?x rdf:type ?c.
    ?x wdt:P161 ?y
}
```

Unger et al. Template-based Question Answering over RDF Data. WWW 2012

Template-based QA

How many {Entity} did {Entity} star in?

How many films did Leonardo DiCaprio star in?


```
SELECT COUNT(?x) WHERE {
    ?x rdf:type wd:Q11424.
    ?x wdt:P161 wd:Q38111
}
```

wd:Q11424	Film
wd:Q38111	Leonardo DiCaprio

Unger et al. Template-based Question Answering over RDF Data. WWW 2012

Pre-defined SPARQL templates

Template-based QA

Pros

- Independent of the KG size
- Fast & parallelizable
- Explainable query results

Cons

- **Manual** curation of templates (100+ is already hard to sustain)
- Each new question formulation will require a **new** template
- **Hard-coded** to the KG schema (ontology)

Unger et al. Template-based Question Answering over RDF Data. WWW 2012

Pipelines: Semantic Parsing

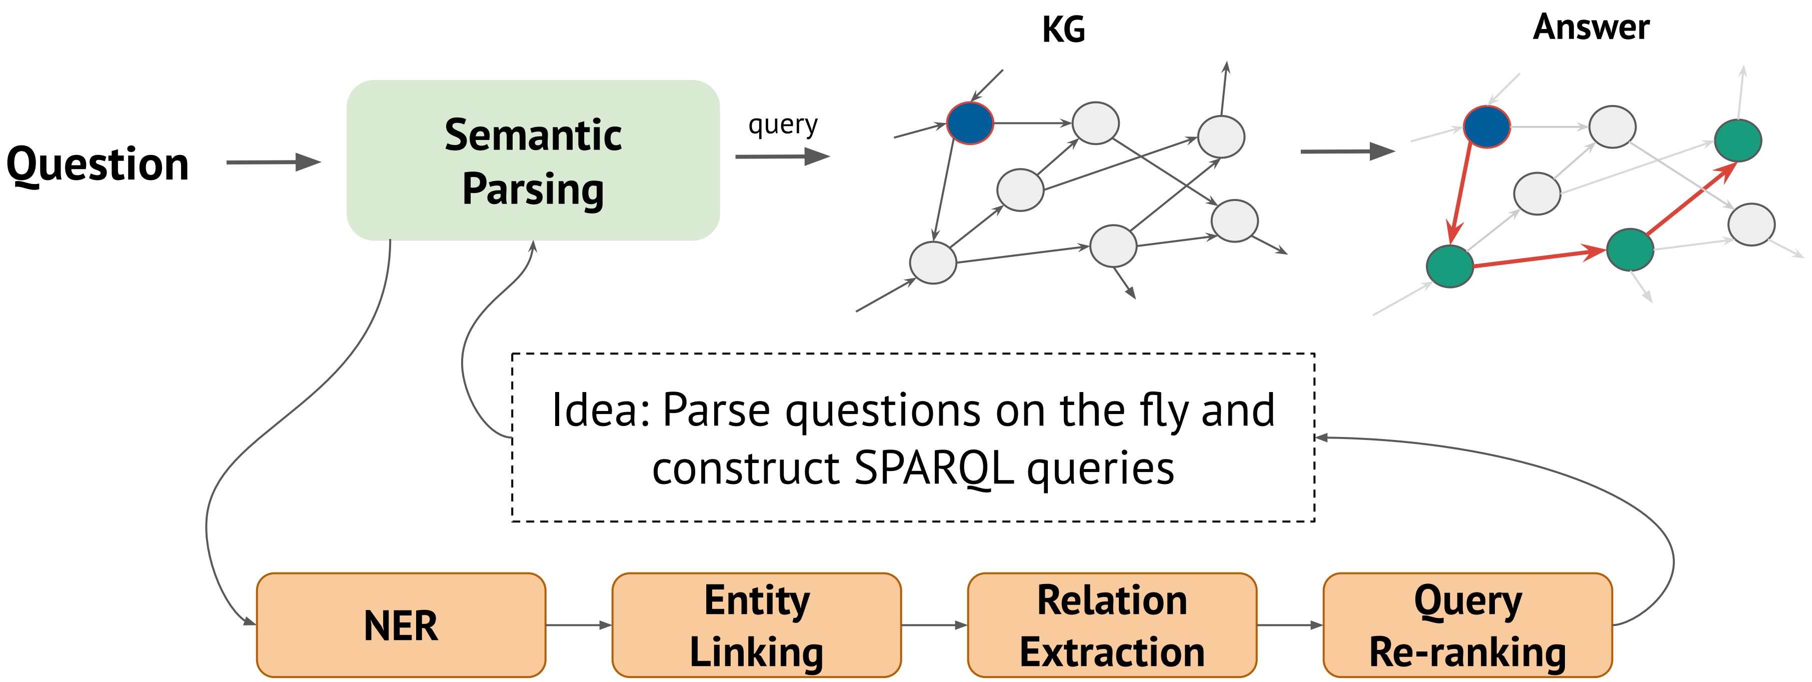

Pipelines: QAmp

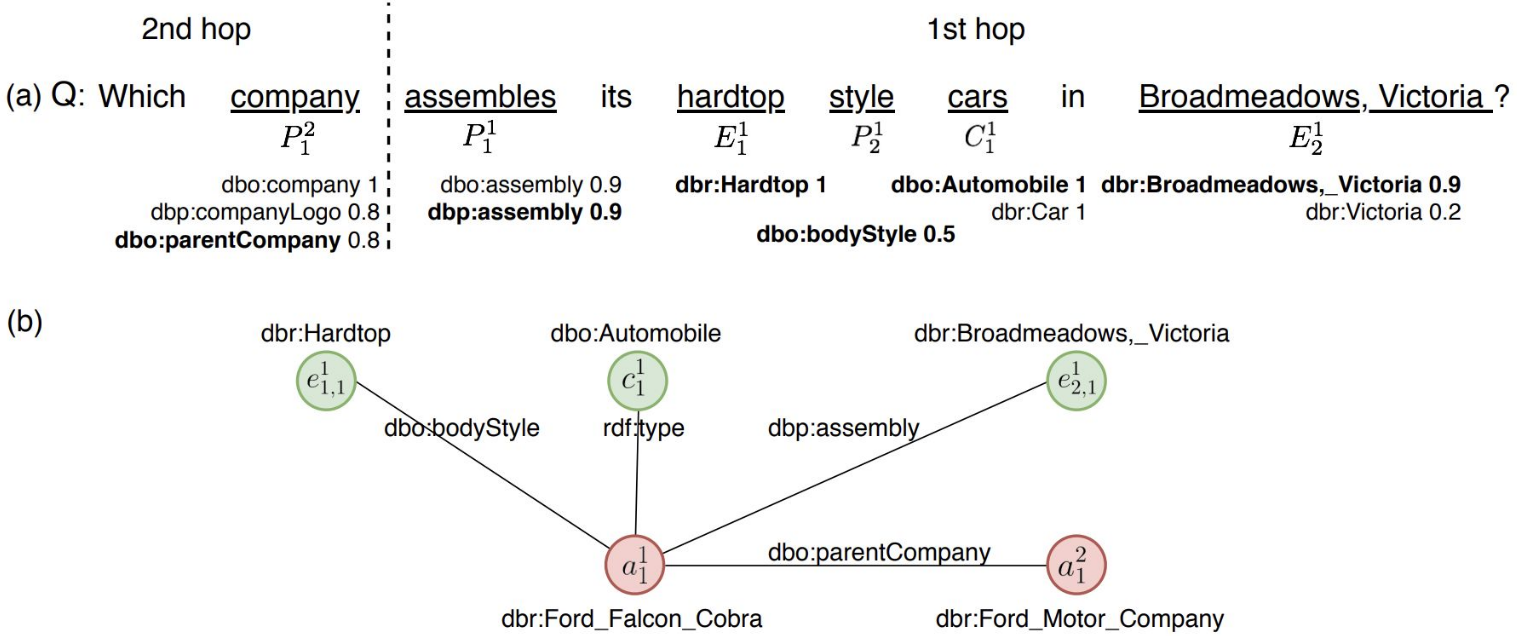

Figure 1: (a) A sample question Q highlighting different components of the question interpretation model: references and matched URIs with the corresponding confidence scores, along with (b) the illustration of a sample KG subgraph relevant to this question. The URIs in bold are the correct matches corresponding to the KG subgraph.

Vakulenko et al. Message Passing for Complex Question Answering over Knowledge Graphs. CIKM 2019

Pipelines: QAmp

- Storage & Querying & subgraph retrieval: HDT
- Entity Linking: ElasticSearch + FastText

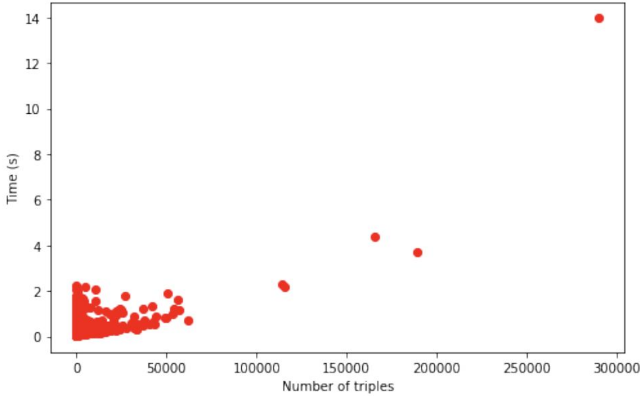

Figure 5: Processing times per question from the LC-QuAD test split (Min: 0.01s Median: 0.68s Mean: 0.72s Max: 13.97s)

Approach	P	R	F	Runtime
WDAqua	0.22*	0.38	0.28	1.50 s/q
QAmp (our approach)	0.25	0.50	0.33	0.72 s/q

Vakulenko et al. Message Passing for Complex Question Answering over Knowledge Graphs. CIKM 2019

Pipelines: Krantikari

Idea: (1) mine core chains (relation paths) from a KG;
(2) re-rank the chains using slot attention

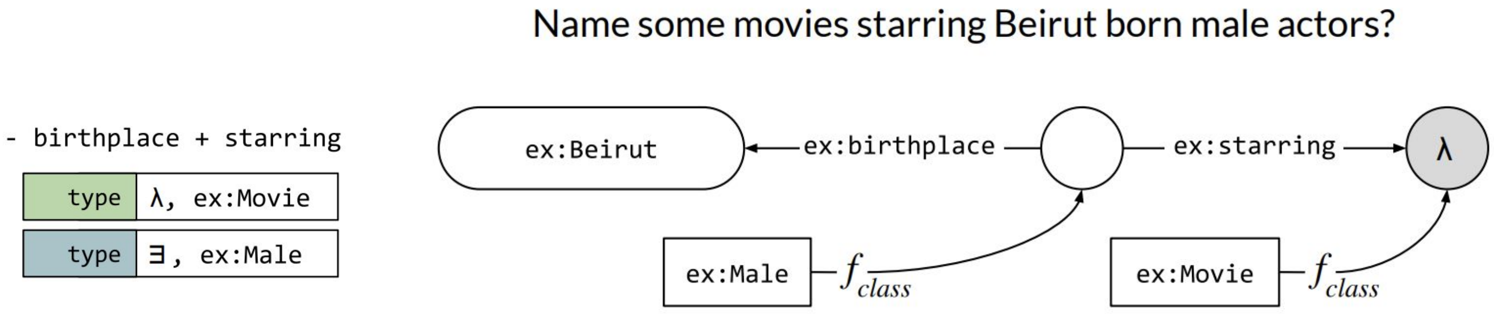

(a) *Question, and corresponding Query Graph*

Maheshwari, Trivedi et al. Learning to Rank Query Graphs for Complex Question Answering over Knowledge Graphs. ISWC 2019

Pipelines: Krantikari

LSTM

	LC-QuAD					QALD-7				
	CCA	MRR	P	R	F1	CCA	MRR	P	R	F1
BiLSTM [9]	0.61	0.70	0.63	0.75	0.68	0.28	0.41	0.20	0.36	0.26
CNN [11]	0.44	0.55	0.49	0.61	0.54	0.31	**0.45**	0.20	0.33	0.25
DAM [16]	0.57	0.66	0.59	0.72	0.65	0.28	0.40	0.20	0.36	0.26
HRM [24]	0.62	0.71	0.64	0.77	0.70	0.28	0.40	0.15	0.31	0.20
Slot-Matching (LSTM)	**0.63**	**0.72**	**0.65**	**0.78**	**0.71**	**0.31**	0.44	**0.28**	**0.44**	**0.34**

Table 1: Performance on LC-Quad and QALD-7. The reported metrics are core chain accuracy (CCA), mean reciprocal rank (MRR) of the core chain rankings, as well as precision (P), recall (R), and the F1 of the execution results of the whole system.

BERT

	QALD-7	LC-QuAD
BERT	**0.23**	0.67
Slot Matching (BERT)	0.18	**0.68**

Table 3: CCA for slot matching model, as proposed in Sec 4.2 initialized with the weights of BERT-Small, compared with regular transformers initialized with the same weights.

Maheshwari, Trivedi et al. Learning to Rank Query Graphs for Complex Question Answering over Knowledge Graphs. ISWC 2019

NL 2 SPARQL

Pros

- Some supervised ML can be applied
- Transfer learning works
- Component ↑ -> Performance ↑
- Explainable query results

Cons

- **Fast** retrieval & communication to the KG is **essential**
- A **Snowball effect** of components error propagation
- **Brute-force** heuristics, e.g., extract a 2-hop subgraph & rank; extract all k-long relation paths and rank

Template-based KGQA

Play around today with DeepPavlov!

```
from deeppavlov import configs, build_model

kbqa_model = build_model(configs.kbqa.kbqa_cq, download=True)
kbqa_model(['Magnus Carlsen is a part of what sport?'])
>>> ["chess"]

kbqa_model = build_model(configs.kbqa.kbqa_cq_rus, download=True)
kbqa_model(['Когда родился Пушкин?'])
>>> ["1799-05-26"]
```

http://docs.deeppavlov.ai/en/master/features/models/kbqa.html

Complex Question Answering

Large KGs (Wikidata-scale)	Smaller KGs
Pre-defined SPARQL templates	Neural Multi-Hop Reasoning
NL -> SPARQL pipelines	Query Embedding

Where did Canadian citizens with Turing Award graduate?

```
SELECT ?y WHERE {
     ?x :win        :TuringAward .
     ?x :citizen    :Canada .
     ?x :graduate   ?y . }
```

query →

Structured Sources

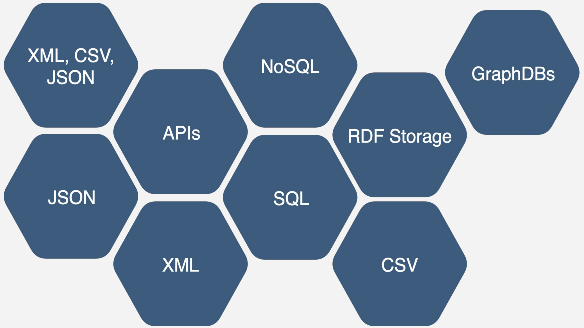

KGs are sparse and incomplete

Neural Multi-Hop Reasoning & Query Embedding

Where did Canadian citizens with Turing Award graduate?

```
SELECT ?y WHERE {
     ?x :win         :TuringAward .
     ?x :citizen     :Canada .
     ?x :graduate    ?y . }
```

embed →

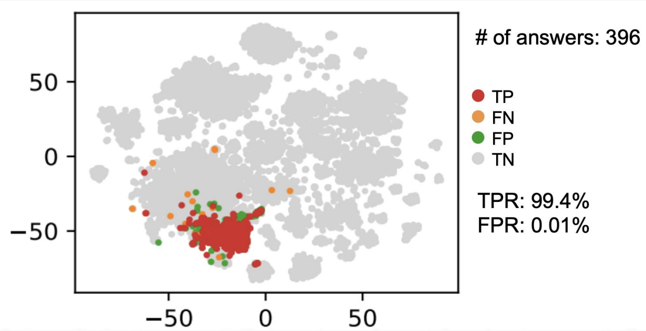

Execution in a vector space

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020
Daza et al. Message Passing Query Embedding. GRL @ ICML 2020
cs224w. snap.stanford.edu

EmbedKGQA

Idea: score a triple (entity, question, answer)

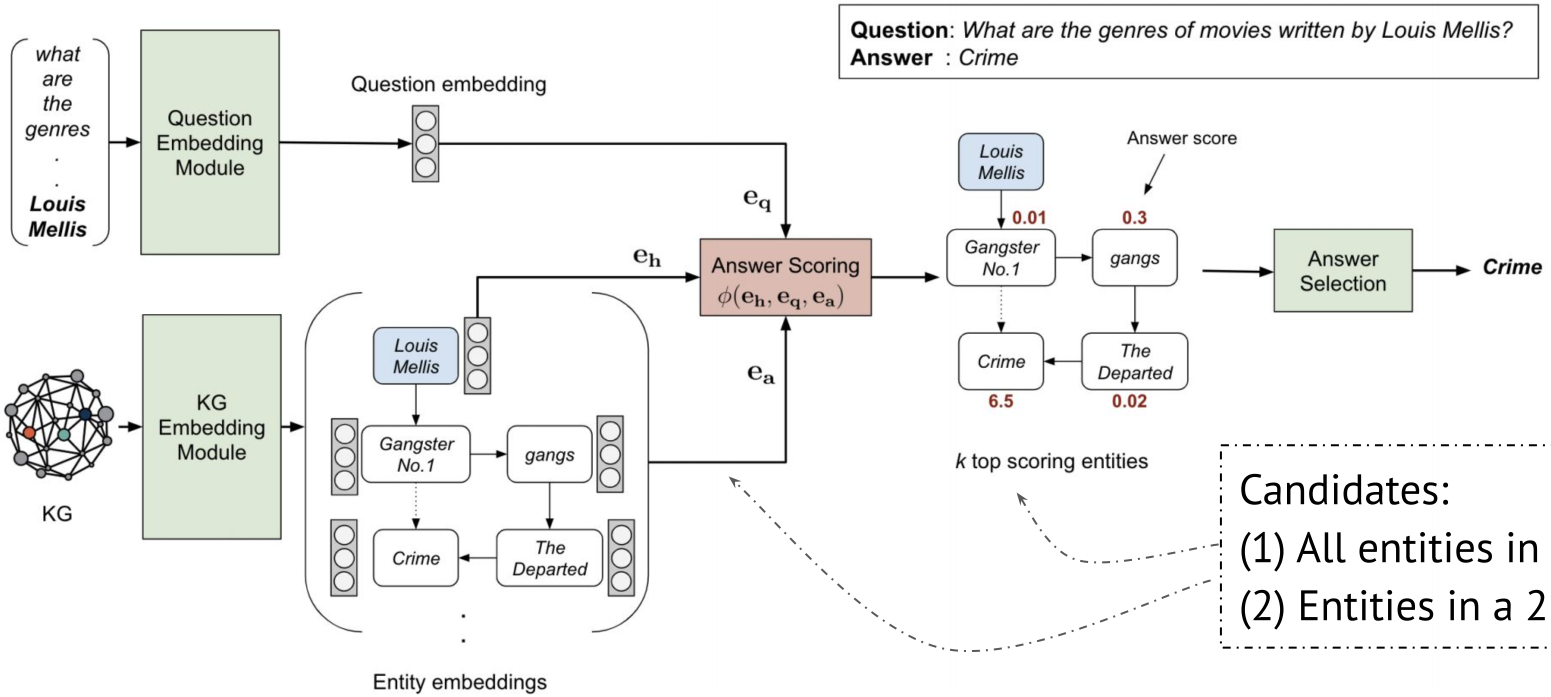

Candidates:

(1) All entities in a KG

(2) Entities in a 2-hop N

Saxena et al. Improving Multi-hop Question Answering over Knowledge Graphs using Knowledge Base Embeddings. ACL 2020

EmbedKGQA

Additional training task: 1-hop link prediction

43k entities
9 relations
135k triples

Model	MetaQA KG-Full			MetaQA KG-50		
	1-hop	2-hop	3-hop	1-hop	2-hop	3-hop
VRN	**97.5**	89.9	62.5	-	-	-
GraftNet	97.0	94.8	77.7	64.0 (91.5)	52.6 (69.5)	59.2 (66.4)
PullNet	97.0	**99.9**	91.4	65.1 (92.4)	52.1 (90.4)	59.7 (85.2)
KV-Mem	96.2	82.7	48.9	63.6 (75.7)	41.8 (48.4)	37.6 (35.2)
EmbedKGQA (Ours)	**97.5**	98.8	**94.8**	**83.9**	**91.8**	**70.3**

1.8M entities
5.7M triples
8 GPUs ;)

Model	WebQSP KG-Full	WebQSP KG-50
KV-Mem	46.7	32.7 (31.6)
GraftNet	66.4	48.2 (49.7)
PullNet	**68.1**	50.1 (51.9)
EmbedKGQA	66.6	**53.2**

Saxena et al. Improving Multi-hop Question Answering over Knowledge Graphs using Knowledge Base Embeddings. ACL 2020

Query2Box

Subset of SPARQL - EPFO queries: Conjunctive + disjunction

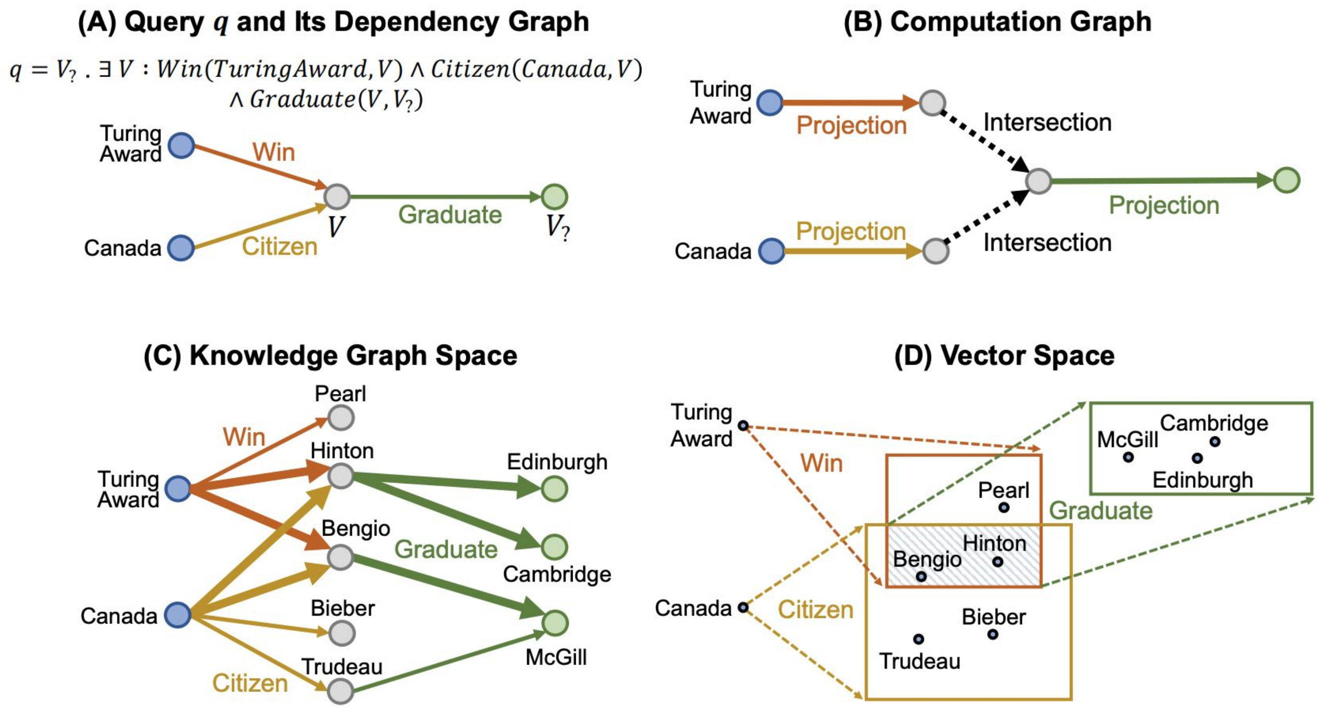

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

Idea: represent entities as d-dimensional boxes!

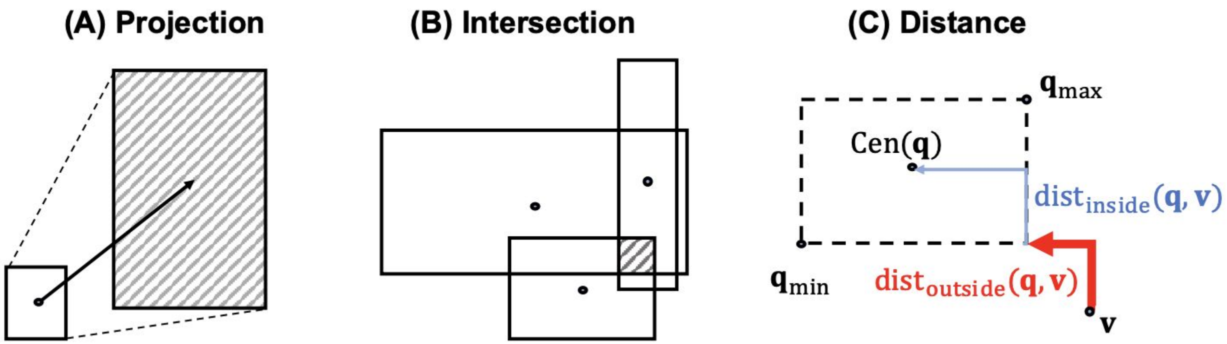

Figure 2: The geometric intuition of the two operations and distance function in QUERY2BOX. **(A)** Projection generates a larger box with a translated center. **(B)** Intersection generates a smaller box lying inside the given set of boxes. **(C)** Distance $\mathrm{dist}_{\mathrm{box}}$ is the weighted sum of $\mathrm{dist}_{\mathrm{outside}}$ and $\mathrm{dist}_{\mathrm{inside}}$, where the latter is weighted less.

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

Query Embedding

Idea: represent entities as d-dimensional boxes!

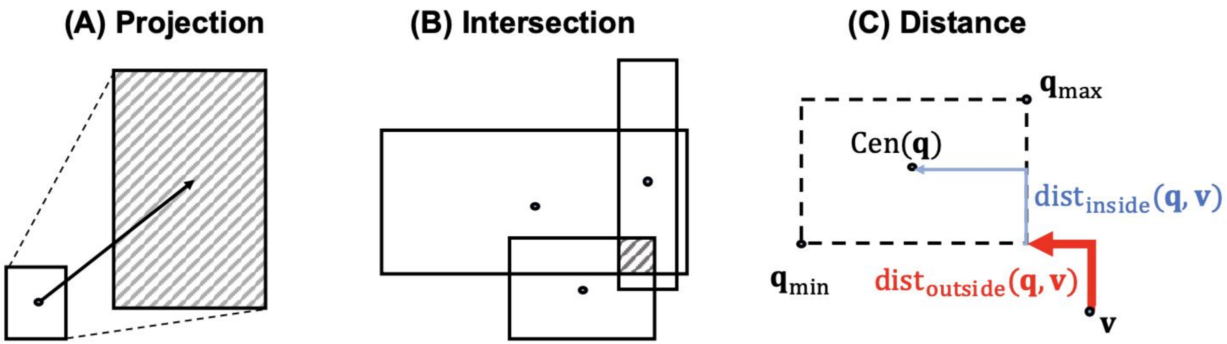

Figure 2: The geometric intuition of the two operations and distance function in QUERY2BOX. **(A)** Projection generates a larger box with a translated center. **(B)** Intersection generates a smaller box lying inside the given set of boxes. **(C)** Distance $\mathrm{dist}_{\mathrm{box}}$ is the weighted sum of $\mathrm{dist}_{\mathrm{outside}}$ and $\mathrm{dist}_{\mathrm{inside}}$, where the latter is weighted less.

Intersection operator

$$\mathrm{Cen}(\mathbf{p}_{\mathrm{inter}}) = \sum_i \mathbf{a}_i \odot \mathrm{Cen}(\mathbf{p_i}), \quad \mathbf{a}_i = \frac{\exp(\mathrm{MLP}(\mathbf{p_i}))}{\sum_j \exp(\mathrm{MLP}(\mathbf{p_j}))},$$

$$\mathrm{Off}(\mathbf{p}_{\mathrm{inter}}) = \mathrm{Min}(\{\mathrm{Off}(\mathbf{p_1}), \dots, \mathrm{Off}(\mathbf{p_n})\}) \odot \sigma(\mathrm{DeepSets}(\{\mathbf{p_1}, \dots, \mathbf{p_n}\})),$$

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

Theorem

Modeling any EPFO query needs O(|E|) params

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

Theorem Modeling any EPFO query needs O(|E|) params

Approach Convert all queries to the DNF (union as the last step)

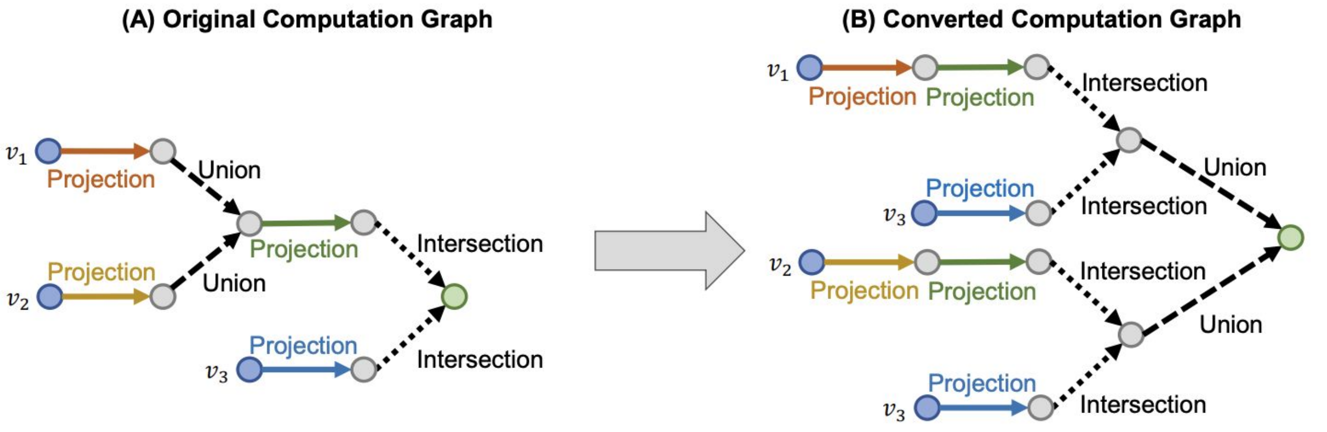

Figure 3: Illustration of converting a computation graph of an EPFO query into an equivalent computation graph of the Disjunctive Normal Form.

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

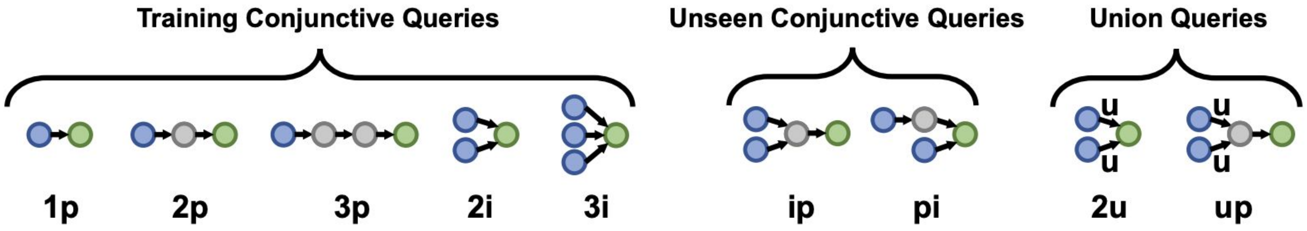

15k entities
1345 relations
500k triples

15k entities
237 relations
300k triples

65k entities
200 relations
120k triples

Method	Avg	1p	2p	3p	2i	3i	ip	pi	2u	up
FB15k										
Q2B	**0.484**	**0.786**	**0.413**	**0.303**	**0.593**	**0.712**	**0.211**	**0.397**	**0.608**	**0.33**
GQE	0.386	0.636	0.345	0.248	0.515	0.624	0.151	0.310	0.376	0.273
GQE-DOUBLE	0.384	0.630	0.346	0.250	0.515	0.611	0.153	0.320	0.362	0.271
FB15k-237										
Q2B	**0.268**	**0.467**	**0.24**	**0.186**	**0.324**	**0.453**	**0.108**	**0.205**	**0.239**	**0.193**
GQE	0.228	0.402	0.213	0.155	0.292	0.406	0.083	0.17	0.169	0.163
GQE-DOUBLE	0.23	0.405	0.213	0.153	0.298	0.411	0.085	0.182	0.167	0.16
NELL995										
Q2B	**0.306**	**0.555**	**0.266**	**0.233**	**0.343**	**0.48**	**0.132**	**0.212**	**0.369**	**0.163**
GQE	0.247	0.418	0.228	0.205	0.316	0.447	0.081	0.186	0.199	0.139
GQE-DOUBLE	0.248	0.417	0.231	0.203	0.318	0.454	0.081	0.188	0.2	0.139

Table 2: H@3 results of QUERY2BOX vs. GQE on FB15k, FB15k-237 and NELL995.

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

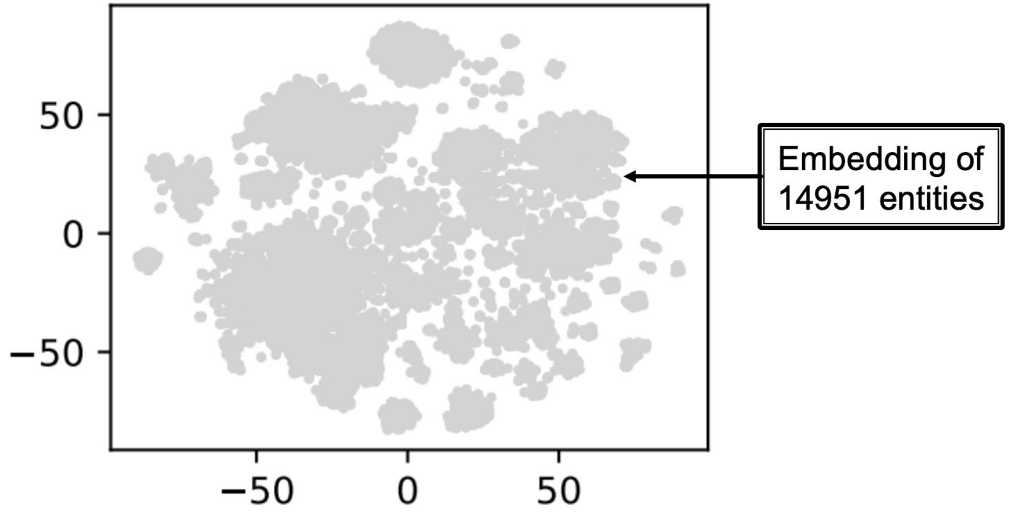

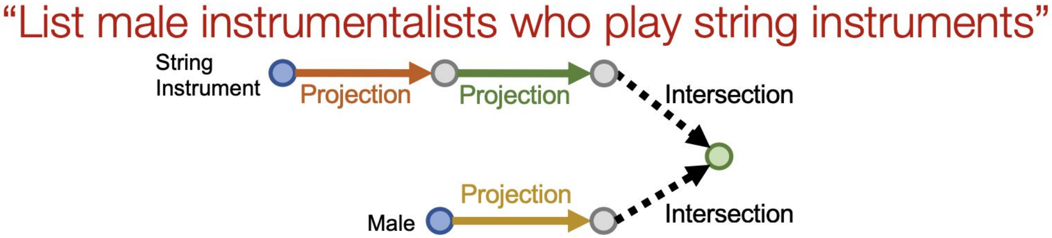

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

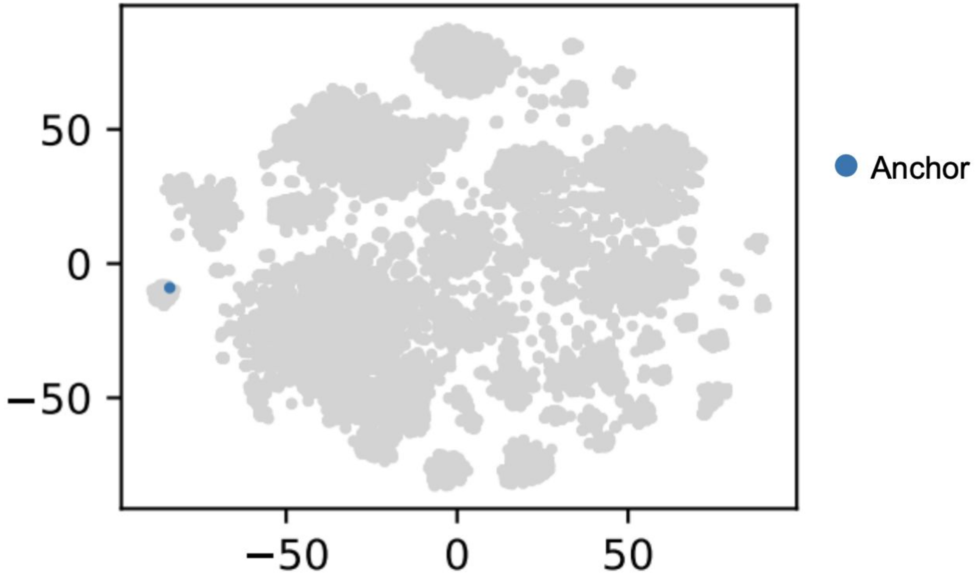

"List male instrumentalists who play string instruments"

String Instrument

Query2Box

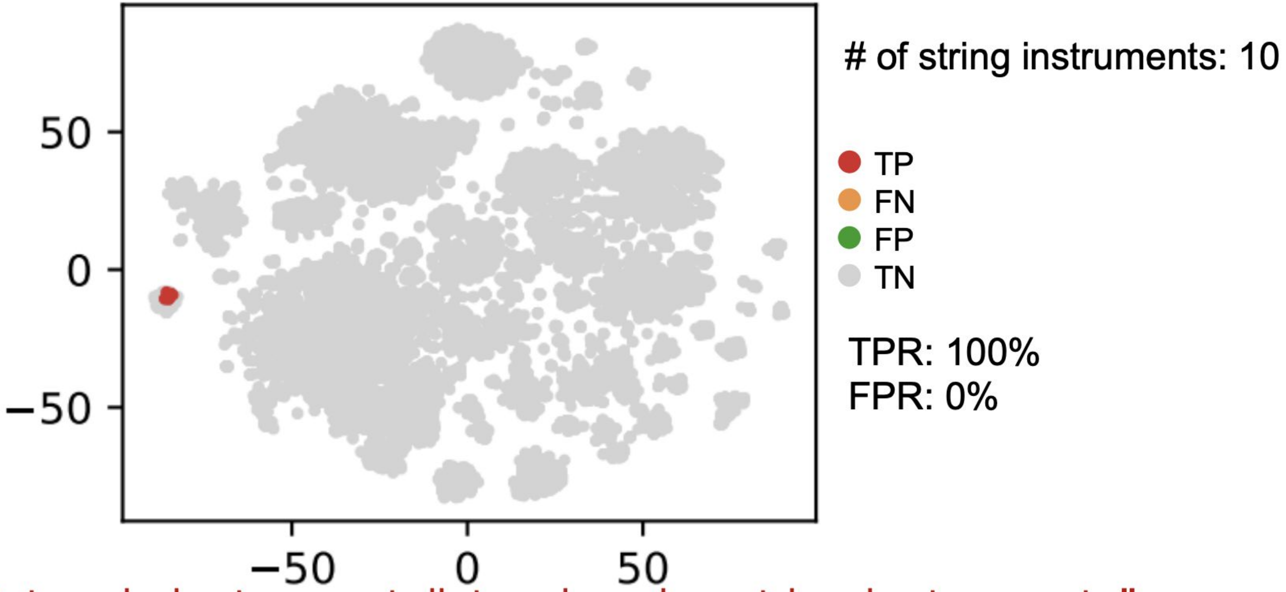

“List male instrumentalists who play <u>string instruments</u>”

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

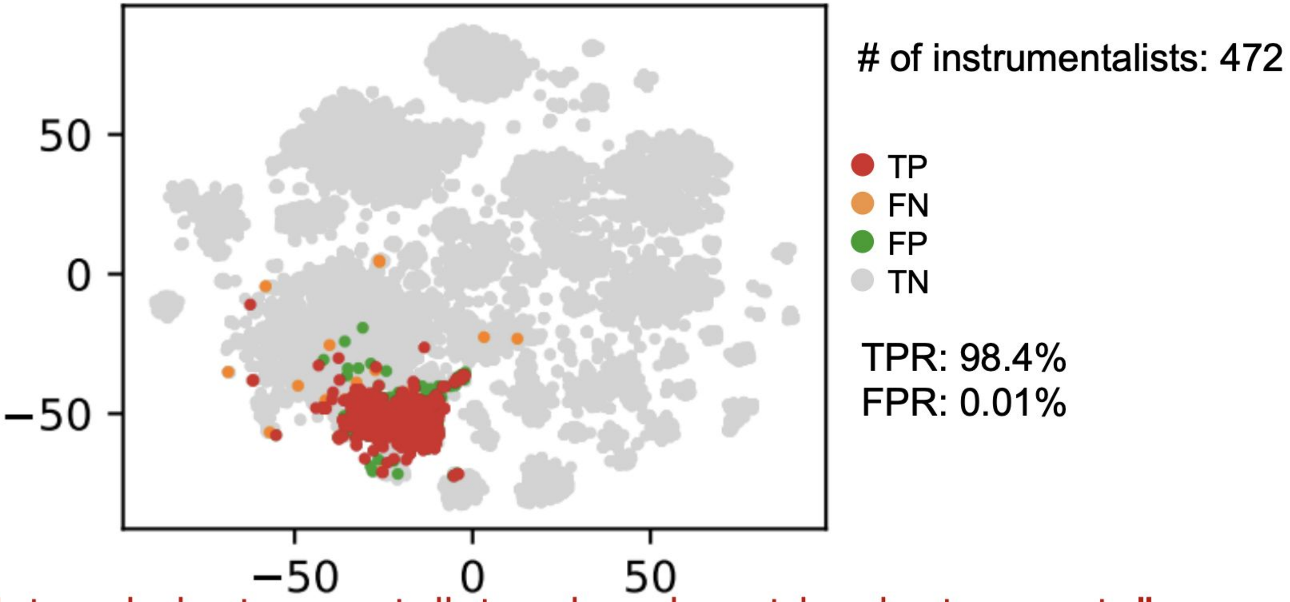

"List male instrumentalists who play string instruments"

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020

Query2Box

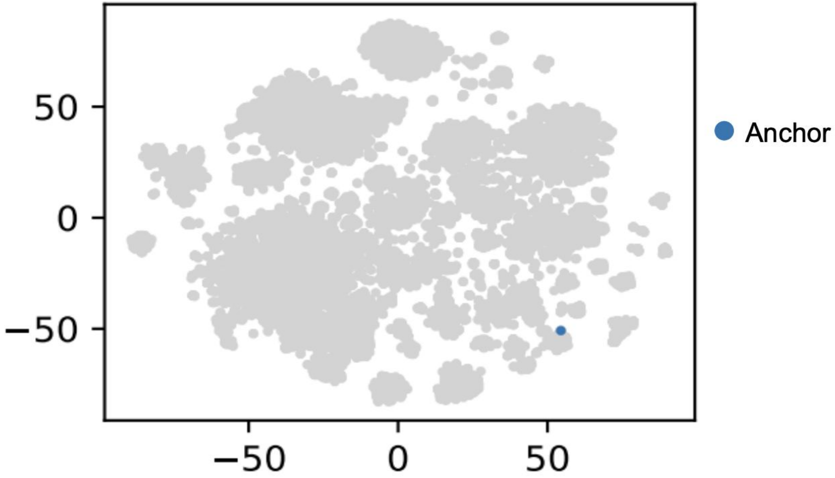

"List male instrumentalists who play string instruments"

Male

Query2Box

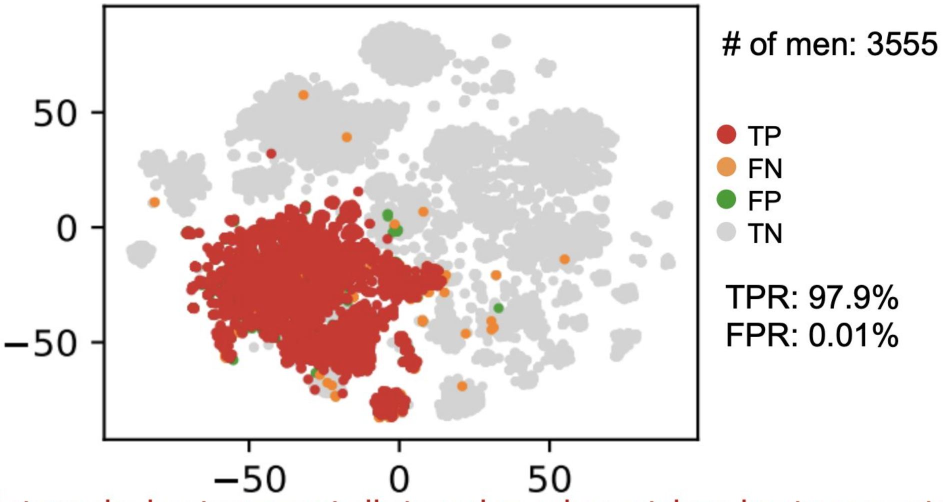

"List male instrumentalists who play string instruments"

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020
cs224w. snap.stanford.edu

Query2Box

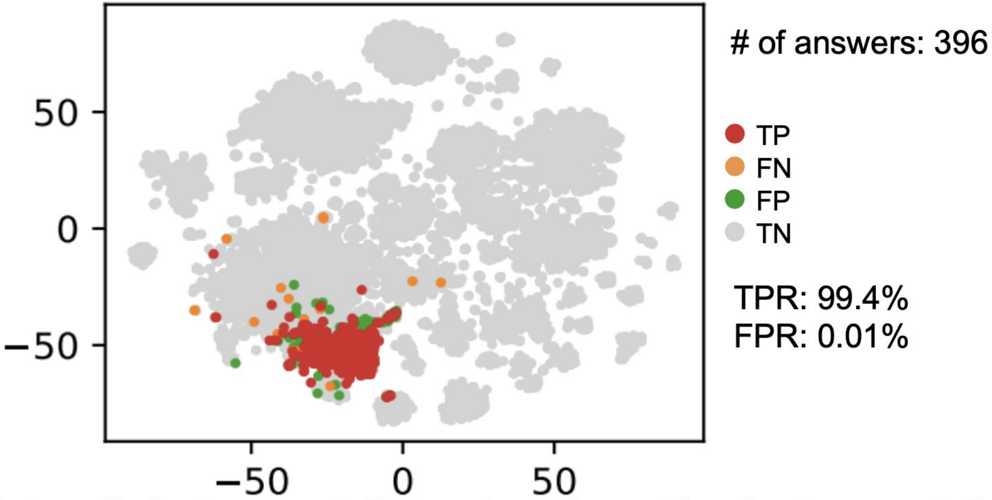

"List male instrumentalists who play string instruments"

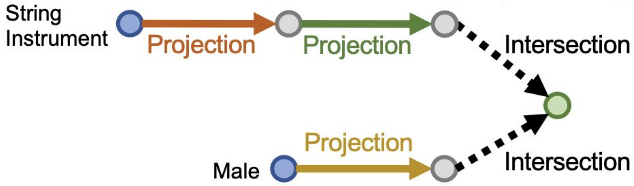

Ren et al. Query2box: Reasoning over Knowledge Graphs in Vector Space Using Box Embeddings. ICLR 2020
cs224w. snap.stanford.edu

Neural Reasoning

Pros

- No database & query engine needed
- No SPARQL / other queries
- We can work with incomplete KGs: infer new facts and predict links

Cons

- **Hardly scalable** to large graphs
- Often **not explainable** results
- Computationally **expensive**
- Problems handling **literals**

Challenges: Answer Verbalization

Input Question

Where was the author of the theory of relativity educated?

Answer

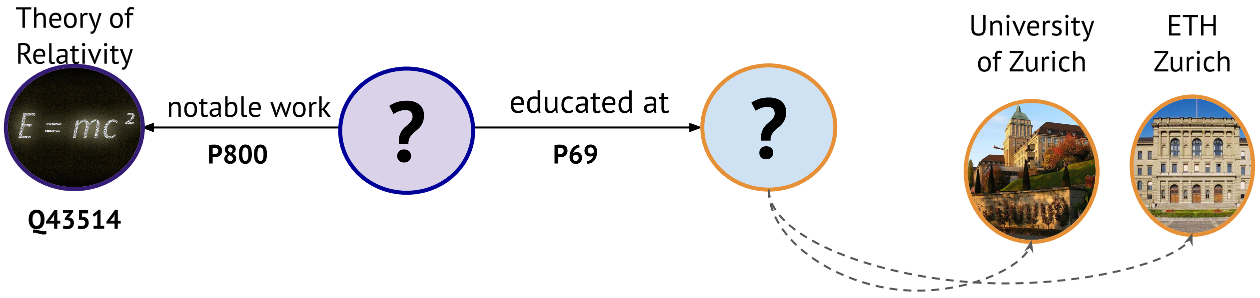

Human-friendly Answer

The author of the theory of relativity educated at the University of Zurich and ETH Zurich.

Kacupaj et al. VQuAnDa: Verbalization QUestion ANswering Dataset. ESWC 2020

Challenges: Hyper-Relational KGs

Where did Albert Einstein receive his degree in physics?

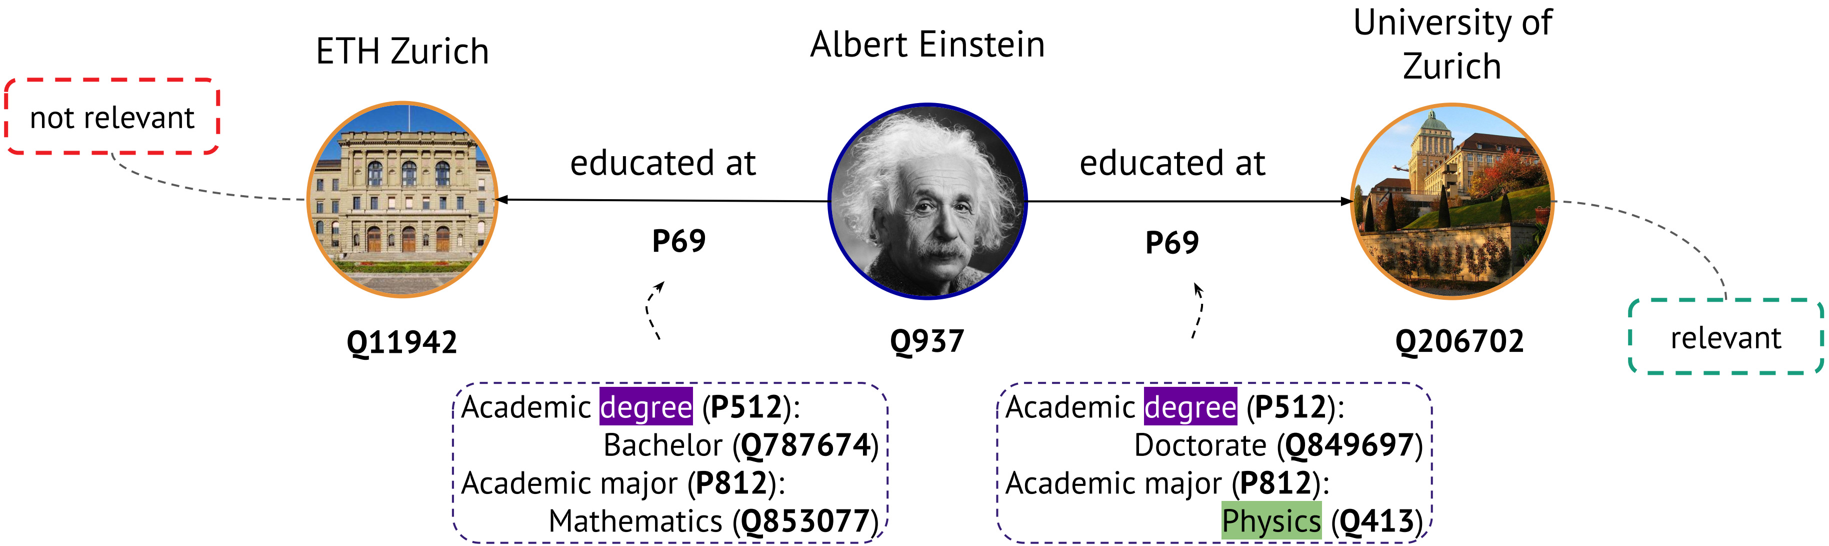

Galkin et al. Message Passing for Hyper-Relational Knowledge Graphs. EMNLP 2020

Multirelational GNN Encoders for KGs

$$\mathbf{h}_v^{(k)} = f\left(\sum_{u \in \mathcal{N}(v)} \mathbf{W}^{(k)} \mathbf{h}_u^{(k-1)}\right)$$

Vanilla GCN [1]: no relations

$$\mathbf{h}_v^{(k)} = f\left(\sum_{(u,r) \in \mathcal{N}(v)} \mathbf{W}_r^{(k)} \mathbf{h}_u^{(k-1)}\right)$$

R-GCN [2]: a whole matrix W per relation

$$\boldsymbol{h}_v = f\left(\sum_{(u,r) \in \mathcal{N}(v)} \boldsymbol{W}_{\lambda(r)} \phi(\boldsymbol{x}_u, \boldsymbol{z}_r)\right)$$

CompGCN [3]: a vector z_r per relation +
composition of (h,r) +
only 3 different W: input/output/self-loop

[1] Kipf et al. Semi-supervised Classification with Graph Convolutional Networks. ICLR 2017
[2] Schlichtkrull et al. Modeling Relational Data with Graph Convolutional Networks. ESWC 2018
[3] Vashishth et al. Composition-Based Multi-Relational Graph Convolutional Networks. ICLR 2020

Embedding Hyper-Relational KGs

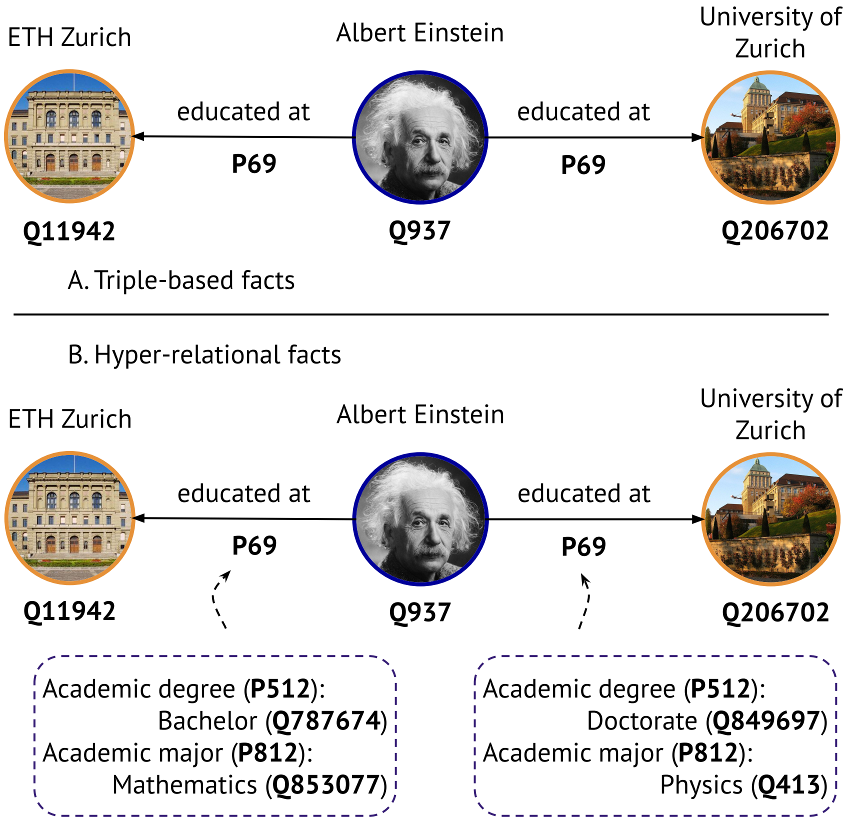

$$\boldsymbol{h}_v = f\left(\sum_{(u,r)\in\mathcal{N}(v)} \boldsymbol{W}_{\lambda(r)} \phi(\boldsymbol{x}_u, \boldsymbol{z}_r) \right)$$

?

- Qualifying relations and entities can be used as main terms in other facts
- Not all facts might have qualifiers

Galkin et al. Message Passing for Hyper-Relational Knowledge Graphs. EMNLP 2020

Embedding Hyper-Relational KGs

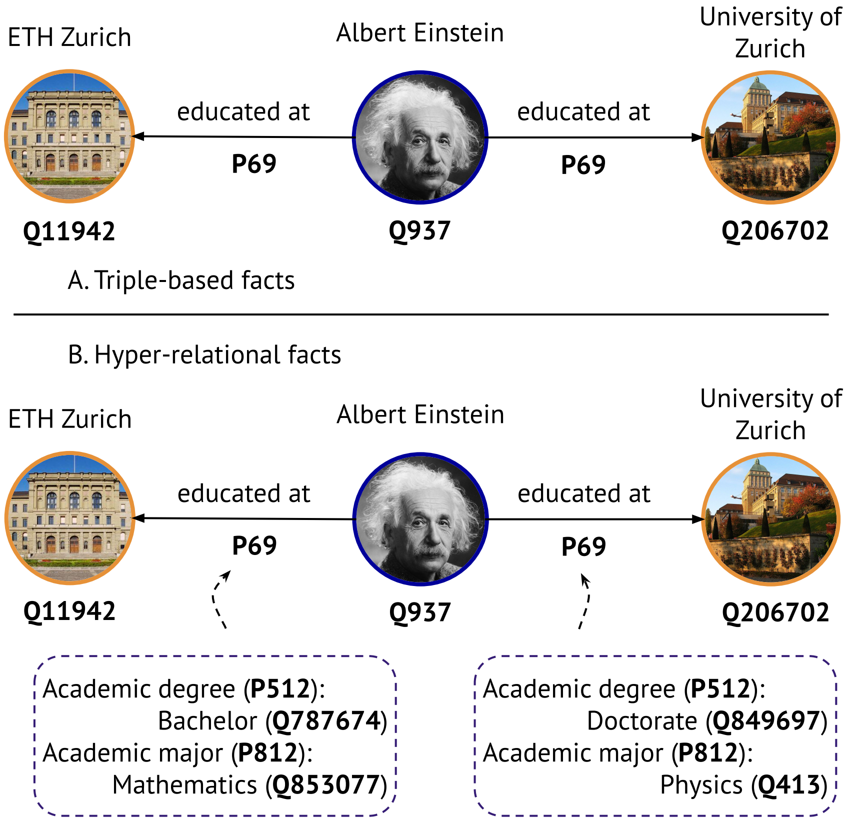

$$\boldsymbol{h}_v = f\left(\sum_{(u,r)\in\mathcal{N}(v)} \boldsymbol{W}_{\lambda(r)} \phi(\boldsymbol{x}_u, \boldsymbol{z}_r) \right)$$

$$\mathbf{h}_v = f\left(\sum_{(u,r)\in\mathcal{N}(v)} \mathbf{W}_{\lambda(r)} \phi_r(\mathbf{h}_u, \gamma(\mathbf{h}_r, \mathbf{h}_q)_{vu}) \right)$$

Galkin et al. Message Passing for Hyper-Relational Knowledge Graphs. EMNLP 2020

★ StarE: Embedding Hyper-Relational KGs

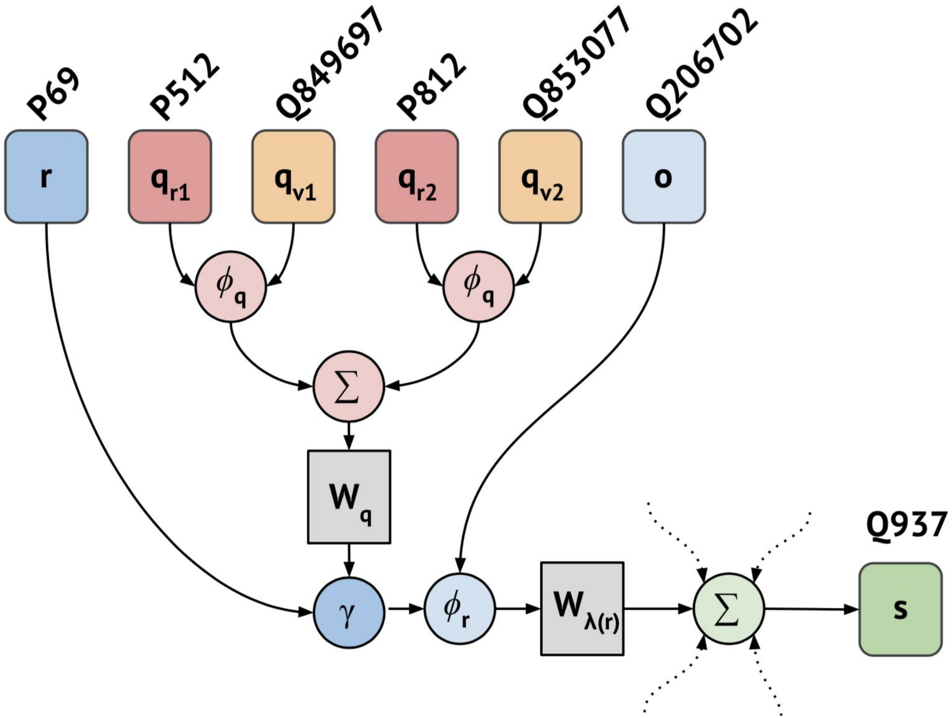

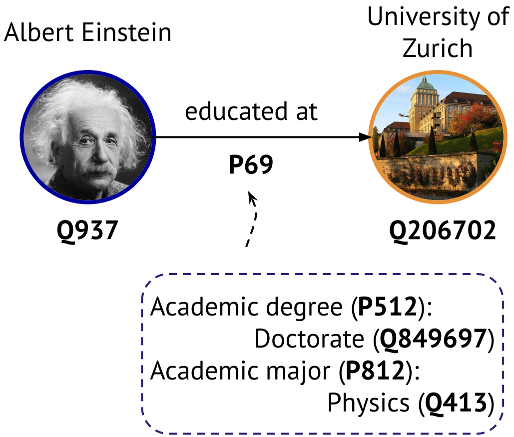

$$\mathbf{h}_v = f\left(\sum_{(u,r)\in\mathcal{N}(v)} \mathbf{W}_{\lambda(r)}\phi_r(\mathbf{h}_u, \gamma(\mathbf{h}_r, \mathbf{h}_q)_{vu})\right)$$

Galkin et al. Message Passing for Hyper-Relational Knowledge Graphs. EMNLP 2020

Hyper-Relational KGs: Link Prediction

Dataset →	WD50K			WD50K (33)			WD50K (66)			WD50K (100)		
Method ↓	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10
Baseline (Transformer (T))	0.275	0.207	0.404	0.218	0.158	0.334	0.270	0.197	0.417	0.351	0.261	0.530
StarE (T) + Transformer(T)	0.308	0.228	0.465	0.246	0.173	0.388	0.297	0.212	0.470	0.380	0.276	0.584
NaLP-Fix	0.177	0.131	0.264	0.204	0.164	0.277	0.334	0.284	0.423	0.458	0.398	0.563
HINGE	0.243	0.176	0.377	0.253	0.190	0.372	0.378	0.307	0.512	0.492	0.417	0.636
Baseline (Transformer (H))	0.286	0.222	0.406	0.276	0.227	0.371	0.404	0.352	0.502	0.562	0.499	0.677
StarE (H) + Transformer(H)	**0.349**	**0.271**	**0.496**	**0.331**	**0.268**	**0.451**	**0.481**	**0.420**	**0.594**	**0.654**	**0.588**	**0.777**

- ★ Hyper-relational models **effectively** leverage qualifiers to improve predictions
- ★ The **more** hyper-relational facts - the **better** are predictions
- ★ The improvement upon triple-only models **grows with the ratio** of hyper-relational edges in the KG

Galkin et al. Message Passing for Hyper-Relational Knowledge Graphs. EMNLP 2020

https://convex.mpi-inf.mpg.de/

Challenges: Dialogue & Sequential QA

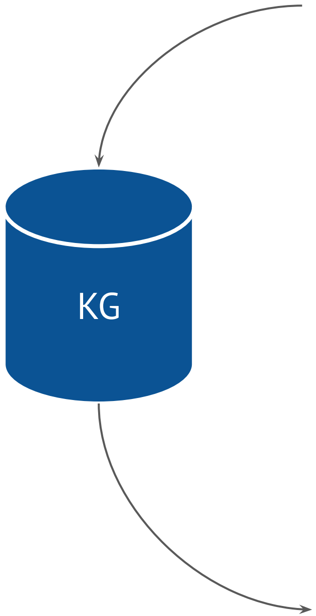

Turn	Books	Movies	Soccer	Music
q^0	*When was the first book of the book series The Dwarves published?*	*Who played the joker in The Dark Knight?*	*Which European team did Diego Costa represent in the year 2018?*	*Led Zeppelin had how many band members?*
a^0	2003	Heath Ledger	Atlético Madrid	4
q^1	*What is the name of the second book?*	*When did he die?*	*Did they win the Super Cup the previous year?*	*Which was released first: Houses of the Holy or Physical Graffiti?*
a^1	The War of the Dwarves	22 January 2008	No	Houses of the Holy
q^2	*Who is the author?*	*Batman actor?*	*Which club was the winner?*	*Is the rain song and immigrant song there?*
a^2	Markus Heitz	Christian Bale	Real Madrid C.F.	No
q^3	*In which city was he born?*	*Director?*	*Which English club did Costa play for before returning to Atlético Madrid?*	*Who wrote those songs?*
a^3	Homburg	Christopher Nolan	Chelsea F.C.	Jimmy Page
q^4	*When was he born?*	*Sequel name?*	*Which stadium is this club's home ground?*	*Name of his previous band?*
a^4	10 October 1971	The Dark Knight Rises	Stamford Bridge	The Yardbirds

Christmann et al. Look before you Hop: Conversational Question Answering over Knowledge Graphs Using Judicious Context Expansion. CIKM 2019

Thanks!

@migalkin
@michael_galkin
@mgalkin
mikhail.galkin@iais.fraunhofer.de
migalkin.github.io