On the Need of Semantics when Tackling Knowledge Graph Refinement under a Machine Learning Perspective

Claudia d'Amato

Computer Science Department
University of Bari "Aldo Moro", Bari, Italy

20th International Conference on Principles of Knowledge Representation and Reasoning (KR 2023)

September 6, 2023

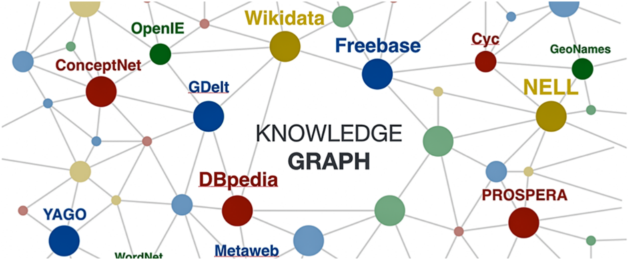

[1]

Open KG
online with content freely accessible

- BabelNet
- DBpedia
- Freebase
- Wikidata
- YAGO
- ....

Enterprise KG
for commercial usage

- Google
- Amazon
- Facebook
- LinkedIn
- Microsoft
- ....

[1] picture from https://www.csee.umbc.edu/courses/graduate/691/fall19/07/

Applications

- e-Commerce
- Semantic Search
- Fact Checking
- Personalization
- Recommendation
- Medical decision support system
- Question Answering
- Machine Translation
- ...

Research Fields

- Information Extraction
- Natural Language Processing
- Machine Learnig (ML)
- Knowledge Representation
- Web
- Robotics
- ...

Knowledge Graph: Definition [Hogan *et al.*, 2021]

A graph of data intended to convey knowledge of the real world

- conforming to a graph-based data model
- nodes represent entities of interest
- edges represent different relations between these entities
- data graph potentially enhanced with schema

KGs: Main Features

- ontologies employed to define and reason about the semantics of nodes and edges
- RDF, RDFS, OWL representation languages largely adopted
- grounded on the Open World Assumption (OWA)
- very large data collections
- suffer of incompleteness and noise
  - since often result from a complex building process

Machine Learning
&
Knowledge Graphs

ML and KGs

Two perspectives:

KG as input to ML

Goal: improving the performance in many learning tasks, e.g.

- Question Answering (QA)
- image classification
- instance disambiguation
- text summarization
- ....

ML as input to KG

Goal: improving the KG itself

- creating new facts
- creating generalizations
- prototyping
- improving the size, coverage, depth and accuracy of KGs $\rightarrow$ reducing their production costs

Reason Why Semantics is Needed

Numeric-based methods

- highly scalable
- consist of series of numbers without any obvious human interpretation
  - no interpretable models provided
  - impact on interpretability, explainability, trustworthiness of results
- no background knowledge and reasoning capabilities generally exploited
  - only factual information is considered

$$\Downarrow$$

- knowledge within KG
  - only partially considered
  - and not always in a fully correct way (negatives)

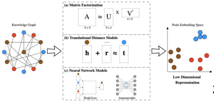

2

DRKG – Drug Repurposing Knowledge Graph

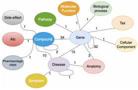

3

---

[2] Picture from D. N. Nicholson et al. Constructing knowledge graphs and their biomedical applications, Computational and Structural Biotechnology Journal, Vol. 18, pp. 1414–1428, (2020) ISSN 2001-0370

[3] Picture from https://github.com/topics/knowledge-graph-embeddings

KG as input to ML

(Infusing KG into Large Language Models)

Large Language Models (LLMs)

LLM: a type AI algorithm that

- uses deep learning techniques and
- massively large data sets
- to understand, summarize, generate and predict new content
  - specifically architected to help generate text-based content
- typically has one billion or more parameters
- is able to understand and generate accurate responses rapidly

Modern LLMs

- use transformer neural networks (transformers)

LLMs take a complex approach that **involves multiple components**

Main Usages of Large Language Models

Once an LLM is trained, a base exists on which performing

- Text generation
- Translation (LLMs trained on multiple languages)
- Content categorization, Content summary and Rewriting
- Conversational AI and chatbots (e.g. ChatGTP)
- Sentiment analysis

LLM: Challenges and Limitations

- Development costs
  - running LLMs require large quantities of expensive graphics processing unit hardware and massive data sets
- Operational costs → can be very high for the hosting organization
- Bias (since trained on unlabeled data)
  - no guarantee that known bias is removed
- **Explainability** → almost missing
- Hallucination
  - when providing an inaccurate response not based on trained data
- Complexity
  - can be particularly complex to troubleshoot
- Glitch tokens
  - maliciously designed prompts that cause an LLM to malfunction

Infusing KG into LLMs

Infusing factual triples of KG into a LLM [Moiseev *et al.*, 2022]

Goal: assess whether LLMs can better internalize knowledge from structured data (KG) or from text on QA tasks

Result: models pre-trained on KG outperform the baseline pre-trained on text sentences containing the same knowledge

Open Challenges

What would be the impact of additionally exploiting KG semantics (e.g. concept hierarchy) and reasoning capabilities (generating also additional fact triples)?

- additionally exploiting KG semantics (e.g. concept hierarchy)
- and reasoning capabilities (generating also additional fact triples)?

$\Downarrow$

- experimental study required
- a new solution for incorporating semantics would be needed

ML as input to KG

(KG Refinement by KG Embedding Methods)

(KG Refinement by Symbol-based Learning Methods)

KG Refinement by KG Embedding Methods:

Injecting Semantics

Incompleteness and **noise**

$\Downarrow$

Knowledge Graph Refinement

- Link Prediction: predicts missing links between entities
  - regarded as a learning to rank problem
- Triple Classification: assesses correctness of a statement wrt a KG
  - regarded as a binary classification problem

Very Large Data Collections

$\Downarrow$

New scalable Machine Learning methods

- grounded on numeric-based approaches
  - KG vector embedding models (KGE) largely investigated [Cai *et al.*, 2018]

ML/KGE for KGs: Issues

- CWA (or LCWA) mostly adopted vs. OWA
- schema level information and reasoning capabilities almost disregarded
- no interpretable models $\Rightarrow$ hard to motivate results

KG Embedding Models...

KGE models convert data graph into an optimal low-dimensional space [Cai *et al.*, 2018]

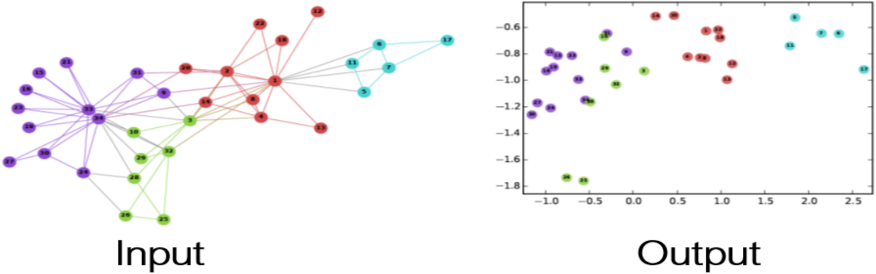

[4]

Graph structural information and properties preserved as much as possible

[4] Picture from https://laptrinhx.com/node2vec-graph-embedding-method-2620064815/

...KG Embedding Models...

KGE methods differ in their main building blocks [Ji *et al.*, 2020]:

- **Representation Space**
  - where represeting relations and entities
  - point-wise, complex, discrete, Gaussian, manifold, etc.
- **Encoding Model**
  - for representing and learning relational interactions
  - linear, factorization, neural models, etc.
- **Scoring Function**
  - for measuring the plausibility of factual triples
  - based on distance, energy, semantic matching, other criteria
- **Auxiliary Information**
  - to be incorporated into the embedding methods
  - text, type, images

...KG Embedding Models

Goal

Learning embeddings s.t.

- score of a valid (positive) triple is higher than
- the score of an invalid (negative) triple

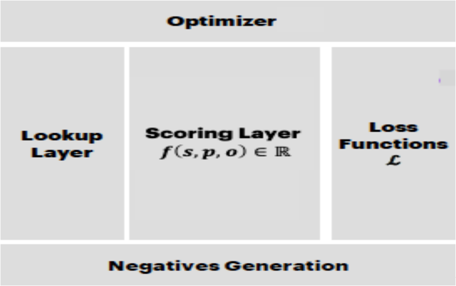

[5]

Negative examples generated by **random corruption of triples**

- false negatives may be generated
- only triple directly observable are considered

[5] Picture from "ECAI-20 Tutorial: Knowledge Graph Embeddings: From Theory to Practice"

Idea: Enhance KGE through Background Knowledge (BK) Injection

[d'Amato *et al.*, 2021c,b]

By two components:

Reasoning: used for generating negative triples

Axioms: `domain`, `range`, `disjointWith`, `functionalProperty`;

BK Injection: defines constraints on functions, corresponding to the considered axioms, guiding the way embedding are learned

Axioms: `equivClass`, `equivProperty`, `inverseOf` and `subClassOf`.

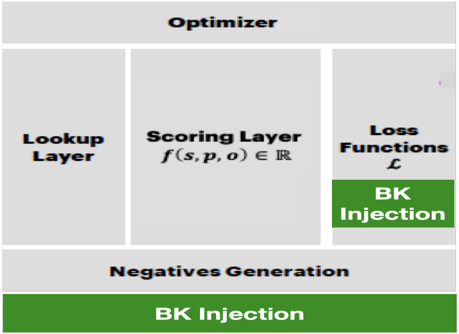

Other KG Embedding Methods Leveraging BK

- Jointly embedding KGs and logical rules [Guo *et al.*, 2016]
  - triples represented as atomic formulae
  - rules represented as complex formulae modeled by t-norm fuzzy logics
- Adversarial training exploiting Datalog clauses encoding assumptions to regularize neural link predictors [Minervini *et al.*, 2017a]

A specific form of BK required, not directly applicable to KGs

An approach to learn embeddings exploiting BK [d'Amato *et al.*, 2021c,b]

TRANSOWL | **TRANSROWL** | **TRANSROWL^R**

TransE
[Bordes *et al.*, 2013]

TransR
[Lin *et al.*, 2015]

Could be applied to more complex KG embedding methods with additional formalization

TransOWL

[d'Amato *et al.*, 2021c]

- Derive further triples to be considered for training via schema axioms
  - `equivClass`, `equivProperty`, `inverseOf` and `subClassOf`
- More complex loss function
  - adding a number of terms consistently with the constraints

$$
\begin{aligned}
L \;=\; & \overbrace{\sum_{\substack{\langle h,r,t\rangle \in \Delta \\ \langle h',r,t'\rangle \in \Delta'}} [\gamma + f_r(h,t) - f_r(h',t')]_+}^{\text{TransE } loss\ \ function} + \sum_{\substack{\langle t,q,h\rangle \in \Delta_{\mathbf{inverseOf}} \\ \langle t',q,h'\rangle \in \Delta'_{\mathbf{inverseOf}}}} [\gamma + f_q(t,h) - f_q(t',h')]_+ \\
& + \sum_{\substack{\langle h,s,t\rangle \in \Delta_{\mathbf{equivProperty}} \\ \langle h',s,t'\rangle \in \Delta'_{\mathbf{equivProperty}}}} [\gamma + f_s(h,t) - f_s(h',t')]_+ + \sum_{\substack{\langle h,\mathbf{typeOf},l\rangle \in \Delta \cup \in \Delta_{\mathbf{equivClass}} \\ \langle h',\mathbf{typeOf},l'\rangle \in \Delta' \cup \Delta'_{\mathbf{equivClass}}}} [\gamma + f_{\mathrm{typeOf}}(h,l) - f_{\mathrm{typeOf}}(h',l')]_+ \\
& + \sum_{\substack{\langle h,\mathbf{subClassOf},p\rangle \in \Delta_{\mathbf{subClass}} \\ \langle h',\mathbf{subClassOf},p'\rangle \in \Delta'_{\mathbf{subClass}}}} [(\gamma - \beta) + f(h,p) - f(h',p')]_+
\end{aligned}
$$

where $q \equiv r^-$, $s \equiv r$ (properties), $l \equiv t$ and $t \sqsubseteq p$ (classes) and $f(h,p) = \|\mathrm{e}_h - \mathrm{e}_p\|$

TransROWL

[d'Amato *et al.*, 2021b]

- TransOWL loss function adopted plus weighting parameters
  - `equivClass`, `equivProperty`, `inverseOf` and `subClassOf`
- TransR score function adopted

$$
\begin{aligned}
L \quad = \quad & \sum_{\substack{\langle h,r,t\rangle \in \Delta \\ \langle h',r,t'\rangle \in \Delta'}} [\gamma + f'_r(h,t) - f'_r(h',t')]_+ + \lambda_1 \sum_{\substack{\langle t,q,h\rangle \in \Delta_{\mathbf{inverseOf}} \\ \langle t',q,h'\rangle \in \Delta_{\mathbf{inverseOf}'}}} [\gamma + f'_q(t,h) - f'_q(t',h')]_+ \\
& + \lambda_2 \sum_{\substack{\langle h,s,t\rangle \in \Delta_{\mathbf{equivProperty}} \\ \langle h',s,t'\rangle \in \Delta_{\mathbf{equivProperty}'}}} [\gamma + f'_s(h,t) - f'_s(h',t')]_+ + \lambda_3 \sum_{\substack{\langle h,\mathbf{typeOf},l\rangle \in \Delta \cup \Delta_{\mathbf{equivClass}} \\ \langle h',\mathbf{typeOf},l'\rangle \in \Delta' \cup \Delta'_{\mathbf{equivClass}}}} [\gamma + f'_{\mathbf{typeOf}}(h,l) - f'_{\mathbf{typeOf}}(h',l')]_+ \\
& + \lambda_4 \sum_{\substack{\langle t,\mathbf{subClassOf},p\rangle \in \Delta_{\mathbf{subClass}} \\ \langle t',\mathbf{subClassOf},p'\rangle \in \Delta_{\mathbf{subClass}'}}} [(\gamma - \beta) + f'(t,p) - f'(t',p')]_+
\end{aligned}
$$

where

- $q \equiv r^-$, $s \equiv r$ (properties), $l \equiv t$ and $t \sqsubseteq p$ (classes)
- the parameters λ_i, $i \in \{1, \ldots, 4\}$, weigh the influence that each function term has during the learning phase

An Alternative Approach: TransROWL^R

[d'Amato *et al.*, 2021c]

Adopting an axiom-based regularization of the loss function as for TransE^R [Minervini *et al.*, 2017b]

- by adding specific constraints to the loss function rather than
- explicitly derive additional triples during training

Loss function

$$
\begin{aligned}
L \quad = \quad & \sum_{\substack{\langle h,r,t\rangle \in \Delta \\ \langle h',r',t'\rangle \in \Delta'}} [\gamma + f'_r(h,t) - f'_r(h',t')]_+ \\
& + \lambda_{\mathbf{1}} \sum_{r \equiv q^- \in \mathcal{T}_{\mathbf{inverseOf}}} \|r + q\| \; + \lambda_{\mathbf{2}} \sum_{r \equiv q^- \in \mathcal{T}_{\mathbf{inverseOf}}} \|M_r - M_q\| \\
& + \lambda_{\mathbf{3}} \sum_{r \equiv p \in \mathcal{T}_{\mathbf{equivProp}}} \|r - p\| \; + \lambda_{\mathbf{4}} \sum_{r \equiv p \in \mathcal{T}_{\mathbf{equivProp}}} \|M_r - M_p\| \\
& + \lambda_{\mathbf{5}} \sum_{e' \equiv e'' \in \mathcal{T}_{\mathbf{equivClass}}} \|e' - e''\| \; + \lambda_{\mathbf{6}} \sum_{s' \subseteq s'' \in \mathcal{T}_{\mathbf{subClass}}} \|\mathbf{1} - \beta - (s' - s'')\|
\end{aligned}
$$

Lesson Learnt from Experiments...

Goal: Assessing the benefit of exploiting BK

- Comparing[6] TransOWL, TransROWL, TransROWL^R over to the original models TransE and TransR as a baseline

Perfomances tested on:

- Link Prediction task
- Triple Classification task
- Standard metrics adopted

KGs adopted:

KG	#Triples	#Entities	#Relationships
DBpedia15K	180000	12800	278
DBpedia100K	600000	100000	321
DBpediaYAGO	290000	88000	316
NELL[7]	150000	68000	272

[6] All methods implemented as publicly available systems https://github.com/Keehl-Mihael/TransROWL-HRS

[7] equivalentClass and equivalentProperty missing; limited number of typeOf-triples; abundance of subClassOf-triples

...Lesson Learnt from Experiments

- Best performance achieved by TransROWL, in most of the cases, and TransROWL^R
- TransROWL slightly superior performance of TransROWL^R
- NELL was aimed at testing in condition of larger incompleteness
  - `equivalentClass` and `equivalentProperty` **missing**
  - low number of typeOf-triples per entity
  - the models showed oscillating performances wrt the baselines

Open Challenges

- further enhance semantic KGE models with additional schema axioms
- extend the framework to more complex KGE models
- improve semantic KGE models for coping with incompleteness $\rightarrow$ further experiments needed

Exploiting Semantics for Providing Explanations to Link Predictions on KGs

A-Posteriori Explanations of Link Predictions...

A-posteriori methods find suitable explanation(s) based on the observed output and the model input, independently on the KGE adopted

Given the predicted triple: $\langle$NickMason, recordLabel, CapitolRecords$\rangle$

why is it provided?

User is able to understand motivations, and trust (or not) the prediction

Example of exmplanation

$\langle$NickMason, associatedBand, PinkFloyd$\rangle$,
$\langle$PinkFloyd, recordLabel, CapitolRecords$\rangle$

Ideally supported by analogous situations to be found in the KG e.g.

$\langle$RingoStarr, recordLabel, Parlophone$\rangle$

for which the computed explanation is:
$\langle$RingoStarr, associatedBand, TheBeatles$\rangle$,
$\langle$TheBeatles, recordLabel, Parlophone$\rangle$.

...A-Posteriori Explanations of Link Predictions

A-posteriori methods: developed solutions

KELPIE [Rossi *et al.*, 2022]: generates necessary and sufficient (path) conditions and an articulated new evaluation protocol

CrossE [Zhang *et al.*, 2019]: embedding model for link predictions providing explanations

- the search for a path linking the subject h and object t of a predicted triple $\langle h, r, t\rangle$
  - **Max lenght 2** $\rightarrow$ six types of paths possible:
    Length 1: $P_1 = \{\langle h, r_s, t\rangle\}$, $P_2 = \{\langle t, r_s, h\rangle\}$
    Length 2: $P_3 = \{\langle e', r_s, h\rangle, \langle e', r', t\rangle\}$, $P_4 = \{\langle e', r_s, h\rangle, \langle t, r', e'\rangle\}$,
    $P_5 = \{\langle h, r_s, e'\rangle, \langle e', r', t\rangle\}$, $P_6 = \{\langle h, r_s, e'\rangle, \langle t, r', e'\rangle\}$,
    where r_s similar to r, r' any other relationship, e' any other entity;
- search driven by similarities between relation/entity embeddings by Euclidean distance
- structural comparisons with other paths in the KG to reinforce the reliability of the explanation found (referred to as support)

Idea: SemanticCrossE provide semantic-based explanations for link predictions on KGs

[d'Amato *et al.*, 2021a]

Extends CrossE by adopting a Semantic Cosine similarity that leads the explanation process

- exploits the underlying KG semantics $\rightarrow$ Domain, Range and Classes considered
- increases the cosine similarity of two (entities / relationships) vector embeddings on the ground of available additional semantic information which is captured by a semantic Score function

Definition (semantic Cosine)

Given KG $\mathcal{K}(\mathcal{E}, \mathbb{R})$, the semantic Cosine measure for two entities $e, e' \in \mathcal{E}$ is defined by:

$$\mathrm{semCos}_{\alpha,\beta}(e, e') = \alpha \cdot \mathrm{sScore}(e, e') + \beta \cdot \mathrm{sim}_{\cos}(\mathbf{e}, \mathbf{e}')$$

where $\mathbf{e}$ the respective entity embedding vector; $\alpha, \beta \in [0, 1]$ chosen s.t. $\alpha + \beta = 1$.

In the case of relations $r, r' \in \mathbb{R}$ the measure is defined analogously.

Definition (semantic Score)

Given $\mathcal{C}$ set of the classes occurring in $\mathcal{K}(\mathcal{E}, \mathbb{R})$, and the functions $Cl\colon \mathcal{E} \to \mathcal{C}$, $Do\colon\ \mathbb{R} \to \mathcal{C}$, and $Ra\colon \mathbb{R} \to \mathcal{C}$ that return, resp., the conjunction of the classes an entity belongs to, and the domain and range of a relation, the semantic Score function for pairs of entities $e, e' \in \mathcal{E}$ is defined by:

$$\mathrm{sScore}(e, e') = \frac{|\mathrm{ret}[Cl(e) \sqcap Cl(e')]|}{|\mathrm{ret}[Cl(e) \sqcup Cl(e')]|}$$

where $\mathrm{ret}_{\mathcal{K}}(C)$ returns the entities that can be proven to belong to a given class C

Analogously, given any two relationships $r, r' \in \mathbb{R}$, it is defined:

$$\mathrm{sScore}(r, r') = \frac{|\mathrm{ret}[Do(r) \sqcap Do(r')]|}{|\mathrm{ret}[Do(r) \sqcup Do(r')]|} + \frac{|\mathrm{ret}[Ra(r) \sqcap Ra(r')]|}{|\mathrm{ret}[Ra(r) \sqcup Ra(r')]|}$$

- Approximated form of semantic Cosine measure (specifically of the semantic Score function) employed [d'Amato *et al.*, 2021a]
- Efficient computation obtained by a preliminary clustering phase [d'Amato *et al.*, 2023]

Lesson Learnt from Experiments...

Goal: Establishing the **impact of an added semantic component when computing explanations** of link prediction results

- Comparing ApproxSemanticCrossE and cosineCrossE to CrossE as baseline
- Code and datasets publicly available[8]

KG	#Triples	#Entities	#Relationships
FB15k-237	310116	14541	237
WN18	151442	40943	18
DBpedia15K	183218	12862	279

- KGs adopted by CrossE considered
- *DBpedia15k* additionally taken for testing the semantic component
- CrossE adopted for the preliminary link prediction phase
- Number k of similar relations; j of most similar entities; $k = j = 3$
- Weights for the approximated semantic cosine similarity: $\alpha = 0.2$; $\beta = 0.8$

[8] https://github.com/pierulohacker/SemanticCrossE/tree/master/explanation

...Lesson Learnt from Experiments

Metrics (as for CrossE):

- ***Recall***: ratio of triples for which the model can generate explanations
- ***Average Support***: number of explanations, on average, for each prediction
- *Recall* and *average support* computed for each of the 6 types of explanation path
- Different ratios (2% and 5%) of predictions considered for building explanations, starting from those ranking higher in the link prediction results

Results:

- ApproxSemanticCrossE showed improved results both in terms of recall and support
- ApproxSemanticCrossE not affected by noisy (irrelevant) explanations as for CrossE and partially cosineCrossE → qualitative evaluation conducted

Open Challenges

- Taking into account additional semantics in KGs (e.g. transitivity, symmetry etc.)
- Injecting BK within KELPIE framework
- Develop a standardized evaluation protocol for explanations

KG Refinement by Symbol-based Methods:

Learning Disjointness Axioms

Symbol-based methods

- usually provide interpretable models
  - e.g. trees, rules, logical formulae, etc.
  - may be exploited for a better understanding of the results
- could be able to exploit a background knowledge
- may be combined with deductive reasoning (e.g to make predictions)
- limited in scalability

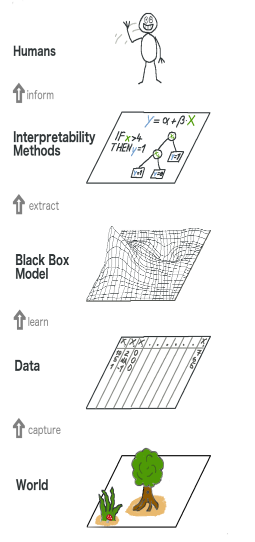

9

[9] Picture from https://jaipancholi.com/model-interpretability

Missing Disjointness Axioms: Issues

Disjointness axioms often missing [Wang *et al.*, 2006]

Problems:

- introduction of noise

$\mathcal{K} = \{ JournalPaper \sqsubseteq Paper,\ ConferencePaper \sqsubseteq Paper,\ ConferencePaper(a), Author(a)\ \}$
$\mathcal{K}$ is Consistent !!!
Cause Axiom: $Author \equiv \neg ConferencePaper$ **missing**

- counterintuitive inferences

$\mathcal{K} = \{ JournalPaper \sqsubseteq Paper,\ ConferencePaper \sqsubseteq Paper,\ ConferencePaper(a)\ \}$

$\mathcal{K} \models JournalPaper(a)$?
Answer: Unknown
Cause Axiom: $JournalPaper \equiv \neg ConferencePaper$ **missing**

- **hard collecting negative examples when adopting numeric approaches**

Observation: extensions of disjoint concepts do not overlap

Question: would it be possible to automatically capture disjointness axioms by analyzing the data configuration/distribution?

Idea: Exploiting **(Conceptual) clustering methods** for mining disjointness axioms

[Rizzo *et al.*, 2021]

Definition (Problem Definition)

Given

- an ontological knowledge base $\mathcal{K} = \langle \mathcal{T}, \mathcal{A} \rangle$
- a set of individuals (aka entities) $\mathsf{I} \subseteq \mathrm{Ind}(\mathcal{A})$

Find

- n pairwise disjoint clusters $\{\mathsf{C}_1, \ldots, \mathsf{C}_n\}$
- for each $i = 1, \ldots, n$, a concept description D_i that describes C_i, such that:
  - $\forall a \in \mathsf{C}_i : \ \mathcal{K} \models D_i(a)$
  - $\forall b \in \mathsf{C}_j, j \neq i : \ \mathcal{K} \models \neg D_i(b)$.
- Hence $\forall D_i, D_j, i \neq j : \ \mathcal{K} \models D_j \sqsubseteq \neg D_i$.

Learning Disjointness Axioms: Developed Methods

Statistical-based approach

- NAR - exploiting negative association rules [Fleischhacker and Völker, 2011]
- PCC - exploiting Pearson's correlation coeff. [Völker *et al.*, 2015]

do not exploit any background knowledge and reasoning capabilities

Disjointness axioms learning/discovery can be hardly performed without symbol-based methods

Terminological Cluster Tree

Defined a method [10] for eliciting disjointness axioms [Rizzo *et al.*, 2021]

- solving a clustering problem via learning Terminological Cluster Trees
- providing a concept description for each cluster

Definition (Terminological cluster tree (TCT))

A binary logical tree where

- a leaf node stands for a cluster of individuals C
- each inner node contains a description D (over the signature of $\mathcal{K}$)
- each departing edge corresponds to positive (left) and negative (right) examples of D

[10] Implemented system publicly available at https://github.com/Giuseppe-Rizzo/TCTnew

Example of TCT

Given $\mathsf{I} \subseteq \mathrm{Ind}(\mathcal{A})$, an example of TCT describing the AI research community

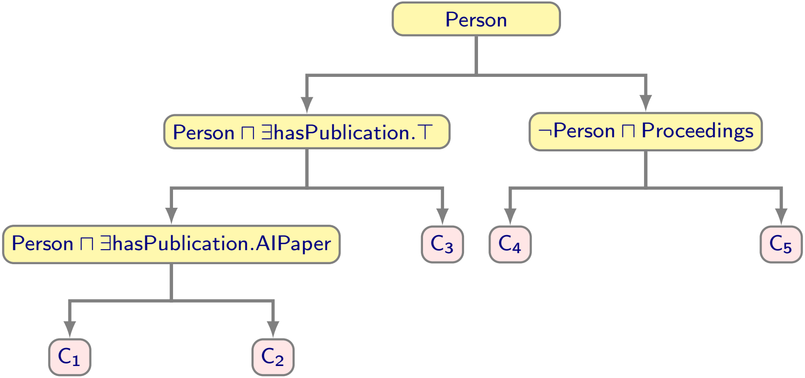

Collecting Disjointness Axioms

Given a TCT T:

Step I:

- Traverse the T to collect the concept descriptions describing the clusters at the leaves
- A set of concepts CS is obtained

Step II:

- A set of candidate axioms A is generated from CS:
  - an axiom $D \sqsubseteq \neg E$ ($D, E \in \mathsf{CS}$) is generated if
    - $D \not\equiv E$ (or $D \not\sqsubseteq E$ or viceversa - reasoner needed)
    - $E \sqsubseteq \neg D$ has not been generated

Inducing a TCT

Given the set of individuals I and $\top$ concept

Divide-and-conquere approach adopted

- **Base Case:** test the stopCondition
  - the cohesion of the cluster I exceeds a threshold ν
    - distance between medoids below a threshold ν
- **Recursive Step** (stopCondition does not hold):
  - a set S of refinements of the current (parent) description C generated
  - the bestConcept $E^* \in$ S is selected and installed as current node
    - the one showing the best cluster separation $\Leftrightarrow$ with max distance between the medoids of its positive P and negative N individuals
  - I is split in:
    - $I_{left} \subseteq$ I $\leftrightarrow$ individuals with the smallest distance wrt the medoid of P
    - $I_{right} \subseteq$ I $\leftrightarrow$ individuals with the smallest distance wrt the medoid of N
    - reasoner employed for collecting P and N

Note: Number of clusters not required - obtained from data distribution

Lesson Learnt from experiments

Experiments performed on ontologies publicly available

- Goal I: Re-discover a target axiom (existing in $\mathcal{K}$)
  - Metrics # discovered axioms and #cases of inconsistency
  - Results:
    - target axioms rediscovered for almost all cases
    - additional disjointness axioms discovered in a significant number
    - limited number of inconsistencies found

Ontology	DL Language	#Concepts	#Roles	#Individuals	#Disj. Ax.s
BioPax	$\mathcal{ALCIF}(D)$	74	70	323	85
NTN	$\mathcal{SHIF}(D)$	47	27	676	40
Financial	$\mathcal{ALCIF}(D)$	60	16	1000	113
GeoSkills	$\mathcal{ALCHOIN}(D)$	596	23	2567	378
Monetary	$\mathcal{ALCHIF}(D)$	323	247	2466	236
DBPedia3.9	$\mathcal{ALCHI}(D)$	251	132	16606	11

Lesson Learnt from experiments [...cont.]

Goal II:

- Re-discover randomly selected target axioms added according to the **Strong Disjointness Assumption** [Schlobach, 2005]
  - two sibling concepts in a subsumption hierarchy considered as disjoint
- comparative analysis with statistical-based methods:
  - PCC - based on Pearson's correlation coefficient [Völker *et al.*, 2015]
  - NAR - exploiting negative association rules [Fleischhacker and Völker, 2011]
- Setting:
  - A copy of each ontology created removing 20%, 50%, 70% of the disjointness axioms
  - **Metrics**: rate of **rediscovered** target axioms, #cases of inconsistency, # addional discovered axioms

Lesson Learnt from experiments [...cont.]

- Results:
  - almost all axioms rediscovered
    - Rate decreases when larger fractions of axioms removed, as expected
  - TCT outperforms PCC and NAR wrt additionally discovered axioms whilst introducing limited inconsistency
    - TCT allows to express complex disjointness axioms
    - PCC and NAR tackle only disjointness between concept names

Exploiting $\mathcal{K}$ as well as the data distribution improves **disjointness axioms discovery**

Open Challenge

Develop an extensive experimental user study on the validity and significance of the complex disjointness axioms discovered by TCT

Conclusions

Conclusions:

- Injecting semantics and exploiting reasoning capabilities may improve the effectiveness of ML solutions for KG
  - Framework for injecting BK into KGE models
  - Solution for injecting semantics when computing a-posteriori explanations to link predictions
- Symbol-based methods useful for supplementing schema level information
  - Conceptual Clustering for Learning Disjointness Axioms

Next Research Challenges:

- Extend the framework for injecting BK to more complex KGE models
- Empower the semantic explanation process with additional schema axioms
- Scalability of symbol-based learning methods still need to be improved
- New solutions required for enhancing LLMs with KG semantics

Thank you


Nicola Fanizzi


Giuseppe Rizzo

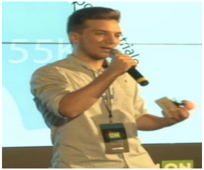
Giovanni Sansaro


Nicola Flavio Quatraro


Pierpaolo Masella


Francesco Benedetti

References

Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O. (2013). Translating embeddings for modeling multi-relational data. In *NIPS 2013*. Curran Assoc., Inc.

Cai, H., Zheng, V. W., and Chang, K. C.-C. (2018). A comprehensive survey of graph embedding: Problems, techniques, and applications. *IEEE Transactions on Knowledge and Data Engineering*, **30**(9), 1616–1637.

d'Amato, C., Masella, P., and Fanizzi, N. (2021a). An approach based on semantic similarity to explaining link predictions on knowledge graphs. In J. He, R. Unland, E. S. Jr., X. Tao, H. Purohit, W. van den Heuvel, J. Yearwood, and J. Cao, editors, *WI-IAT '21: IEEE/WIC/ACM International Conference on Web Intelligence, Melbourne VIC Australia, December 14 - 17, 2021*, pages 170–177. ACM.

d'Amato, C., Quatraro, N. F., and Fanizzi, N. (2021b). Embedding models for knowledge graphs induced by clusters of relations and background knowledge. In N. Katzouris and A. Artikis, editors, *Inductive Logic Programming - 30th International Conference, ILP 2021, Virtual Event, October 25-27, 2021, Proceedings*, volume 13191 of *Lecture Notes in Computer Science*, pages 1–16. Springer.

d'Amato, C., Quatraro, N. F., and Fanizzi, N. (2021c). Injecting background knowledge into embedding models for predictive tasks on knowledge graphs. In R. Verborgh, K. Hose, H. Paulheim, P. Champin, M. Maleshkova, Ó. Corcho, P. Ristoski, and M. Alam, editors, *The Semantic Web - 18th International Conference, ESWC 2021, Virtual Event, June 6-10, 2021, Proceedings*, volume 12731 of *Lecture Notes in Computer Science*, pages 441–457. Springer.

References [...cont.]

d'Amato, C., Benedetti, F., and Fanizzi, N. (2023). Efficient explanation of predictions on dl knowledge graphs through an enhanced similarity search. In O. Kutz, A. Ozaki, and C. Lutz, editors, *DL 2023: 36th International Workshop on Description Logics*. CEUR.

Fleischhacker, D. and Völker, J. (2011). Inductive learning of disjointness axioms. In R. Meersman and et. al., editors, *On the Move to Meaningful Internet Systems: OTM 2011 - Confederated International Conferences: CoopIS, DOA-SVI, and ODBASE 2011, Proceedings, Part II*, volume 7045 of *Lecture Notes in Computer Science*, pages 680–697. Springer.

Guo, S., Wang, Q., Wang, L., Wang, B., and Guo, L. (2016). Jointly embedding knowledge graphs and logical rules. In J. Su, X. Carreras, and K. Duh, editors, *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 192–202. The Association for Computational Linguistics.

Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutierrez, C., Kirrane, S., Gayo, J. E. L., Navigli, R., Neumaier, S., Ngomo, A. N., Polleres, A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., and Zimmermann, A. (2021). *Knowledge Graphs*. Synthesis Lectures on Data, Semantics, and Knowledge. Morgan & Claypool Publishers.

Ji, S., Pan, S., Cambria, E., Marttinen, P., and Yu, P. S. (2020). A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, **33**, 494–514.

References [. . . cont.]

Lin, Y., Liu, Z., Sun, M., Liu, Y., and Zhu, X. (2015). Learning entity and relation embeddings for knowledge graph completion. In *AAAI 2015 Proceedings*, pages 2181–2187. AAAI Press.

Minervini, P., Demeester, T., Rocktäschel, T., and Riedel, S. (2017a). Adversarial sets for regularising neural link predictors. In G. Elidan, K. Kersting, and A. Ihler, editors, *Proceedings of the Thirty-Third Conference on Uncertainty in Artificial Intelligence, UAI 2017, Sydney, Australia, August 11-15, 2017*. AUAI Press.

Minervini, P., Costabello, L., Muñoz, E., Novácek, V., and Vandenbussche, P. (2017b). Regularizing knowledge graph embeddings via equivalence and inversion axioms. In M. Ceci, J. Hollmén, L. Todorovski, C. Vens, and S. Dzeroski, editors, *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2017, Skopje, Macedonia, September 18-22, 2017, Proceedings, Part I*, volume 10534 of *Lecture Notes in Computer Science*, pages 668–683. Springer.

Moiseev, F., Dong, Z., Alfonseca, E., and Jaggi, M. (2022). SKILL: Structured knowledge infusion for large language models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1581–1588, Seattle, United States. Association for Computational Linguistics.

Rizzo, G., d'Amato, C., and Fanizzi, N. (2021). An unsupervised approach to disjointness learning based on terminological cluster trees. *Semantic Web*, **12**(3), 423–447.

References [...cont.]

Rossi, A., Firmani, D., Merialdo, P., and Teofili, T. (2022). Explaining link prediction systems based on knowledge graph embeddings. In Z. G. Ives, A. Bonifati, and A. E. Abbadi, editors, *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*, pages 2062–2075. ACM.

Schlobach, S. (2005). Debugging and semantic clarification by pinpointing. In A. Gómez-Pérez and J. Euzenat, editors, *The Semantic Web: Research and Applications, Second European Semantic Web Conference, ESWC 2005, Heraklion, Crete, Greece, May 29 - June 1, 2005, Proceedings*, volume 3532 of *Lecture Notes in Computer Science*, pages 226–240. Springer.

Völker, J., Fleischhacker, D., and Stuckenschmidt, H. (2015). Automatic acquisition of class disjointness. *Journal of Web Semantics*, **35**, 124–139.

Wang, T. D., Parsia, B., and Hendler, J. A. (2006). A survey of the web ontology landscape. In I. F. Cruz, S. Decker, D. Allemang, C. Preist, D. Schwabe, P. Mika, M. Uschold, and L. Aroyo, editors, *The Semantic Web - ISWC 2006, 5th International Semantic Web Conference, ISWC 2006, Athens, GA, USA, November 5-9, 2006, Proceedings*, volume 4273 of *Lecture Notes in Computer Science*, pages 682–694. Springer.

Zhang, W., Paudel, B., Zhang, W., Bernstein, A., and Chen, H. (2019). Interaction embeddings for prediction and explanation in knowledge graphs. In *Proceedings of WSDM 2019*, pages 96–104. ACM.