Neural-Symbolic Link Prediction Methods, Benchmarks and Explanations for Knowledge Graph Refinement

Claudia d'Amato

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

OLYMPIA 2026 AI Data Infrastructure Summit

November 4 - 5, 2026 - Heraklion, Crete, Greece

Coupling AI Agents with Knowledge Graphs for

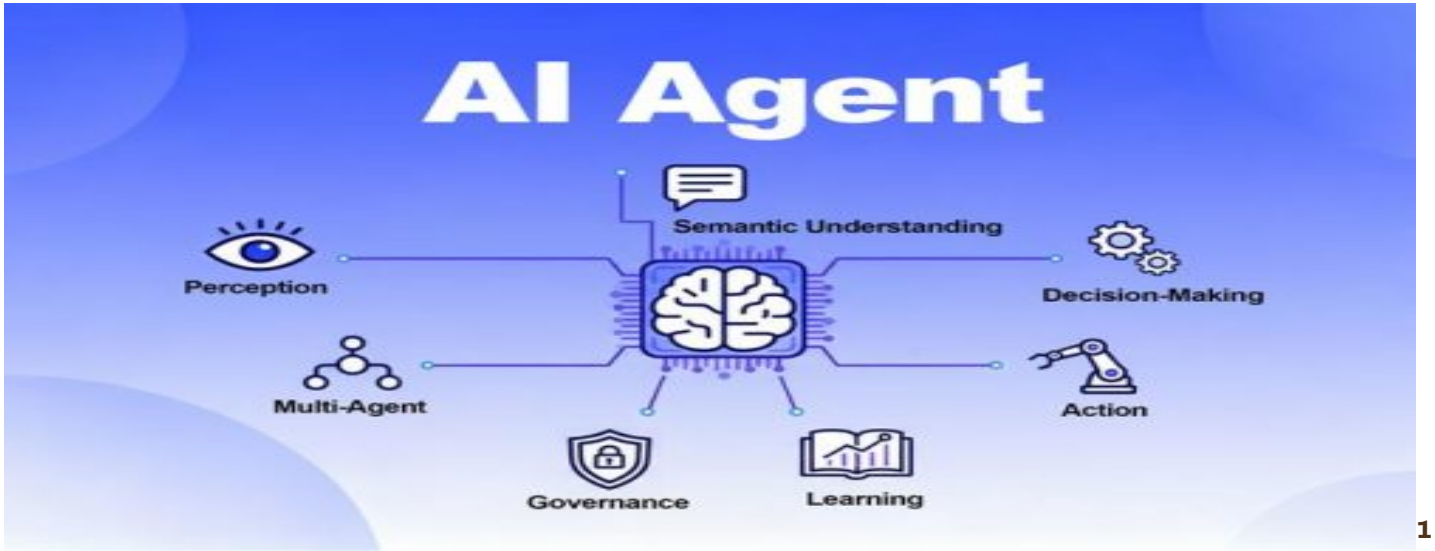

[1]

- reducing hallucinations
- improving multi-hop reaoning
- providing contextual knowledge for improving precision of aswers

[1] picture from https://www.oneas1a.com/what-are-ai-agents/

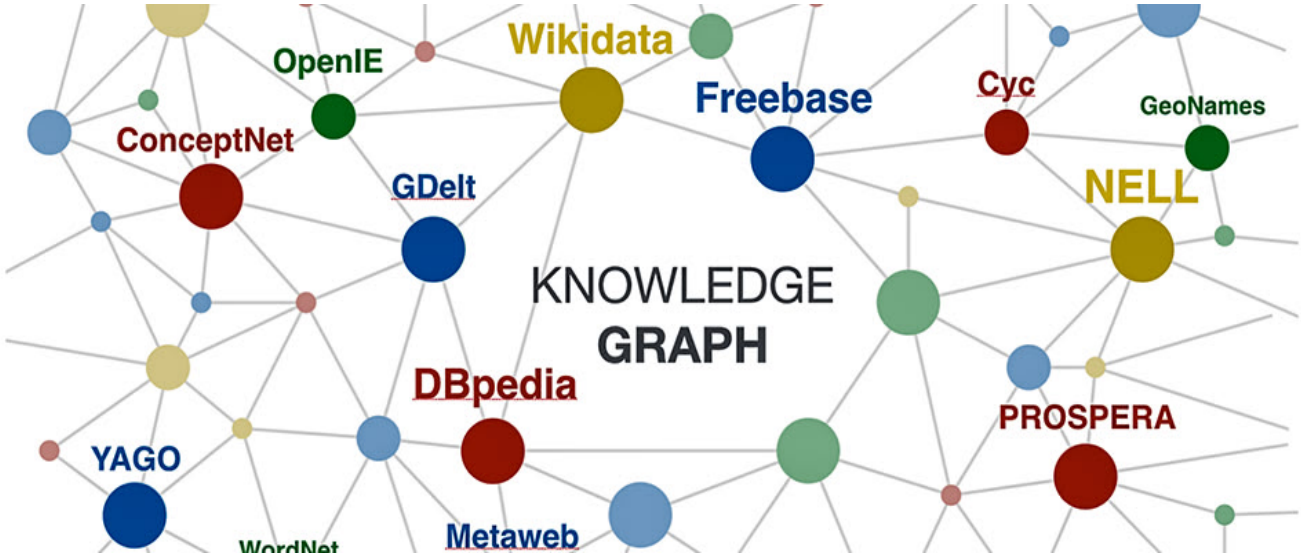

[2]

Open KG
with content freely accessible

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

Enterprise KG
for commercial usage

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

Applications

- e-Commerce
- Semantic Search
- Fact Checking
- Decision support system
- Question Answering
- Machine Translation
- ...

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

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

KG Embedding Models for KG Refinement

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



Input

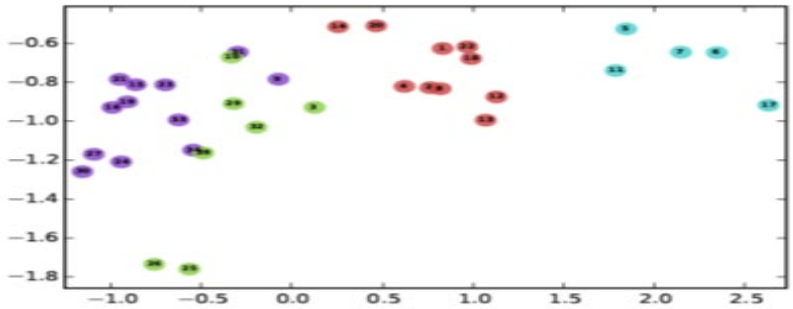

Output

[3]

Graph structural information and properties preserved as much as possible

ML/KGE for KGs: Issues [d'Amato *et al.*2023]

- CWA (or LCWA) mostly adopted vs. OWA
- partial knowledge considered - only triple directly observable
- schema level information and reasoning capabilities almost disregarded
- black box models ⇒ hard to motivate results

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

KG Refinement by NeSy KGE Models:

Injecting Semantics

Enhancing KGE by Injecting Background Knowledge (BK) [d'Amato *et al.*2021b,a] [4] [5]

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`.

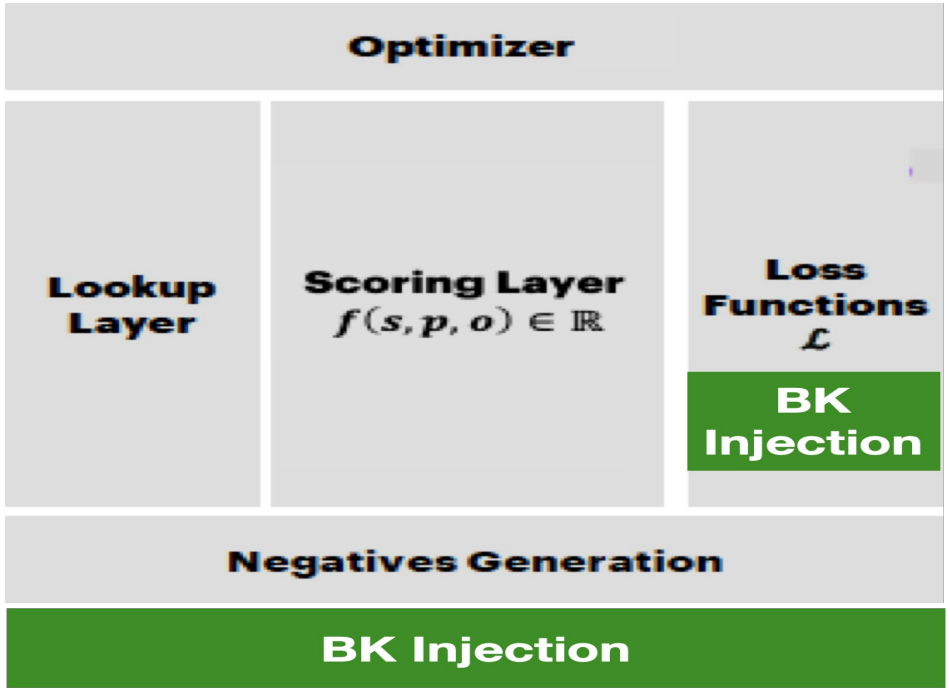

[4] C. d'Amato, N. F. Quatraro, N. Fanizzi: Injecting Background Knowledge into Embedding Models for Predictive Tasks on Knowledge Graphs. ESWC 2021: 441-457 (2021)

[5] C. d'Amato, N. F. Quatraro, N. Fanizzi: Embedding Models for Knowledge Graphs Induced by Clusters of Relations and Background Knowledge. IJCLR 2021 Proceedings (2021)

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.*2017]

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

- BoxEL KGE model the logical structure of ABox and TBox axioms in $\mathcal{EL}^{++}$ Description Logics [Xiong *et al.*2022]
- BoxLitE extend BoxEL focussing on faithful embeddings for DL-Lite$^{\mathcal{H}}$ offering only a proof of concept [Lourenço *et al.*2026]

TransOWL, TransROWL [d'Amato *et al.*2021b] [6]

- 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

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

where $q \equiv r^-$, $s \equiv r$ (properties), $l \equiv t$ and $t \sqsubseteq p$ (classes) and $f(h,p) = \|e_h - e_p\|$

[6] C. d'Amato, N. F. Quatraro, N. Fanizzi: Injecting Background Knowledge into Embedding Models for Predictive Tasks on Knowledge Graphs. ESWC 2021: 441-457 (2021)

Lesson Learnt from Experiments

Goal: Assessing the benefit of exploiting BK

- Comparing[7] TransOWL, TransROWL, over to the original models TransE and TransR as a baseline

KGs adopted:

KG	#Triples	#Entities	#Relationships
DBpedia15K	180000	12800	278
DBpedia100K	600000	100000	321
DBpediaYAGO	290000	88000	316
NELL[8]	150000	68000	272

- **Outcomes:**
  - TransROWL best performing method, in most of the cases

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

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

Evaluating KG Refinement NeSy Solutions:

Need for Proper Benchmarks

Current KGE Evaluation on Link Prediction Tasks

- KGE trained and tested on samples of KGs publicly available
  - e.g. DB100K, YAGO3-10 derived from DBpedia and YAGO3
- Only triple directly observable (data) considered / schema disregarded
- NeSy methods integrate schema with purely data-driven approaches
  - Benchmarks including both data and schema needed [d'Amato *et al.*2023] [9]

[9] C. d'Amato, L. Mahon, P. Monnin, G. Stamou: Machine Learning and Knowledge Graphs: Existing Gaps and Future Research Challenges. TGDK 1(1): 8:1-8:35 (2023)

On Benchmarks Including Data and Schema

A few datasets with both data and schema released [Robert *et al.*2025]

- Obtained by attaching schema fragments (with desired expressivity) to existing data fragments, e.g DB100K
- Heuristic approach adopted, tailored to a specific NeSy method
  - e.g different expressiveness (and generally limited)
  - severely hampers reproducible comparisons among different methods

Proposals differ in [d'Amato *et al.*2021b]:

- how relevant schema fragments are selected
  - generally not grounded in proper ontology modularization methods
  - leading to different schema fragments for datasets from the same KG
- whether datasets are enriched with knowledge derived from reasoning
- the partitioning of datasets into components
  - many datasets released as monoliths, mixing different constructs
- formats - adopt implementation-specific formats
  - rather than standard KG serializations

Benchmarks Including Data and Schema: Requirements

- **Preserving the full original KG (schema) expressiveness**
  - also for NeSy KGR methods leveraging only a few constructs
    - as discarded ones can still be useful for debugging phases
- **Correcting logical flaws, before usage**
  - NeSy methods generally assume flawless knowledge
  - flaws lead to trivial reasoning
- **Enriched with knowledge derived from reasoning**
- **Grounded on modularization**
  - ensuring the adoption of the same schema fragment from the same KG
- **Adopting standard KG serializations**

JediKG Resource [Diliso *et al.*2026][11]

The first resource [10] that systematically addresses the construction of datasets with both schema and data from large and expressive KGs

- Complies the prescribed requirements
- Made by two components
  - **JDeX: a general, semi-automated, interactive and configurable workflow for deriving schema-enhanced datasets on-demand**
    - takes as input any KG in standard Semantic Web (SW) languages
    - derives a dataset in common SW formats
    - includes dataset pre-processing for ML workflows - e.g. to be used ins PyKEEN
  - **JDSet: a curated suite of 11 datasets obtained by applying JDeX to 6 different KGs**, ensuring
    - completeness: covering both data and schema preserving original expressivity
    - portability: usable by both learning and reasoning software
    - serialization in common SW formats
    - release in multiple components concerning different constructs and purposes

[10] Available on Zenodo https://doi.org/10.5281/zenodo.17817931

[11] I. Diliso, R. Barile, N. Fanizzi, C. d'Amato: JediKG: Return of the Schema Building Complete Datasets for Machine Learning and Reasoning on Knowledge Graphs ISWC Proc. (2026)

JDeX workflow

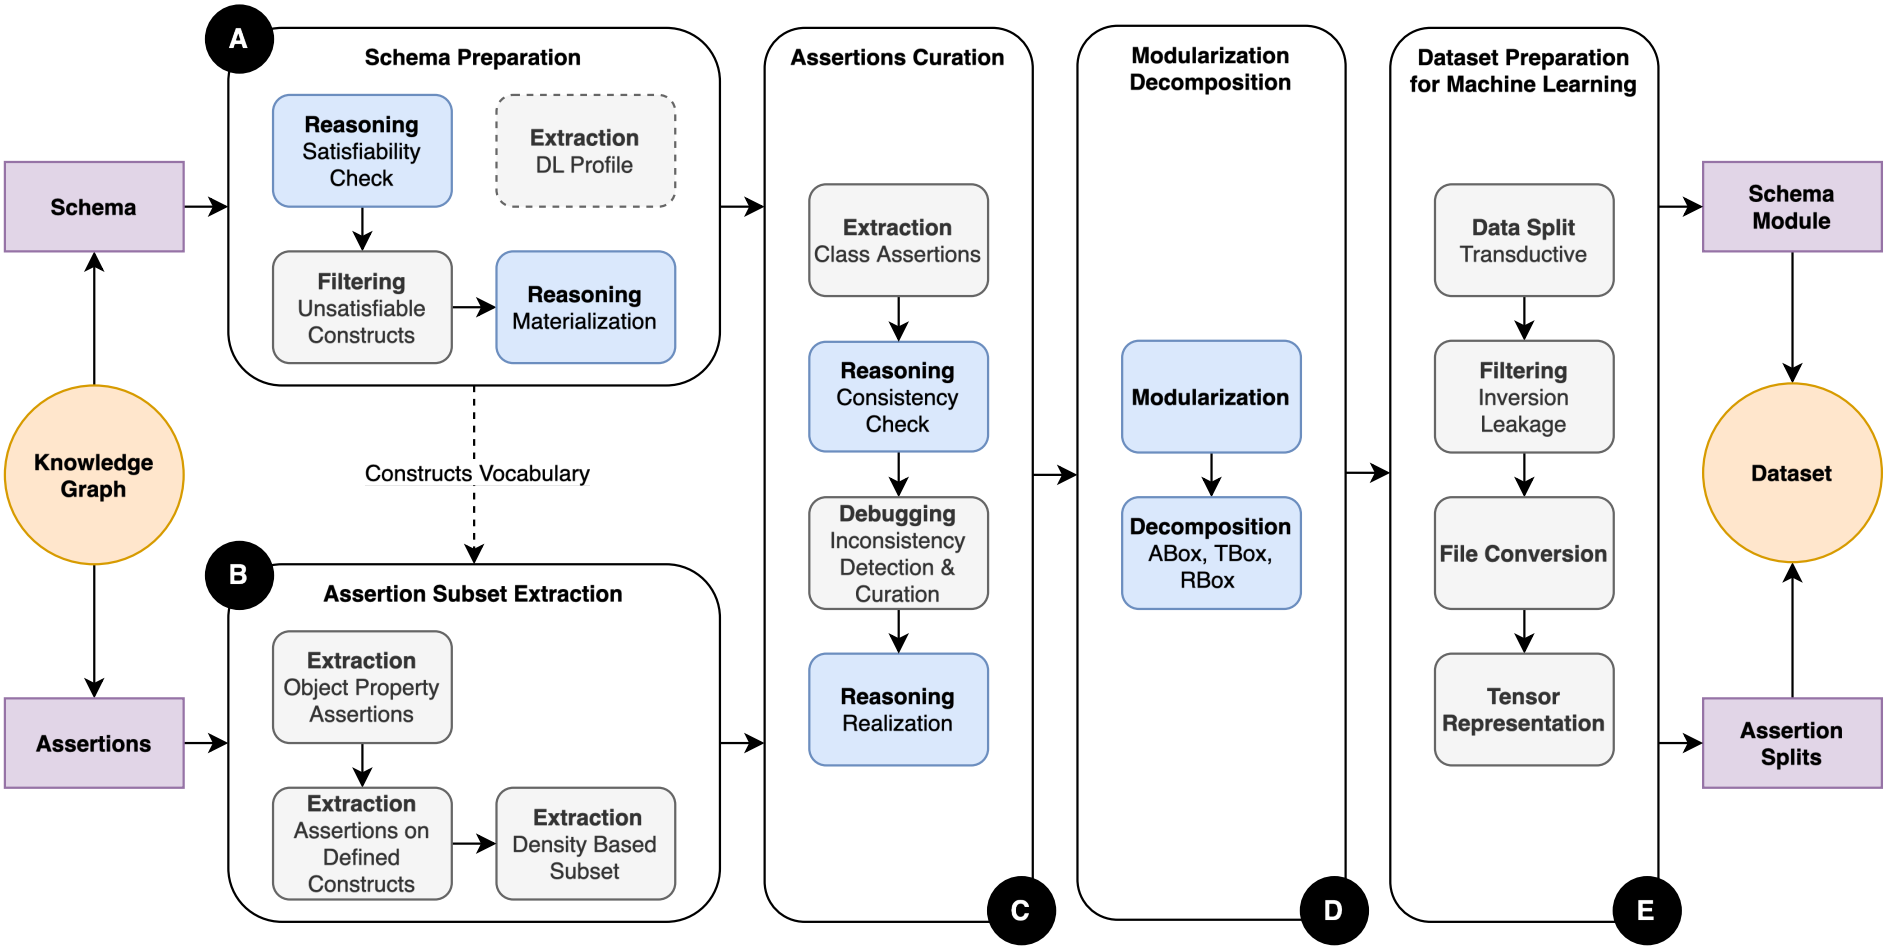

Figure: JDeX workflow phases

JDSet: Dataset Suite

Table 1. Schema axiom type coverage of different datasets. ✗ denotes axioms available in the KG but not in the dataset.

	SOTA							JDSET						
OWL feature	DB100K	DB100K+	YAGO3-10	YAGO3-10+	YAGO3-39K	YAGO4-19	YAGO4-20	DB100K-C	YAGO3-10-C	YAGO3-39K-C	YAGO4-20-C	WHOW-5	ARCO-5	ERA-90
ClassAssertion	✓	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
SubClassOf	✗	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
EquivalentClasses	–	–	–	–	–	–	–	–	–	–	–	–	✓	–
DisjointClasses	✗	✗	–	–	–	✗	✗	✓	–	–	✓	✓	✓	✓
UnionOf	–	–	–	–	–	✗	✗	–	–	–	✓	✓	✓	✓
IntersectionOf	–	–	–	–	–	–	–	–	–	–	–	–	–	–
ComplementOf	–	–	–	–	–	–	–	–	–	–	–	–	✓	–
SomeValuesFrom	–	–	–	–	–	–	–	–	–	–	–	✓	✓	–
AllValuesFrom	–	–	–	–	–	–	–	–	–	–	–	✓	✓	✓
CardinalityRestriction	–	–	–	–	–	–	–	–	–	–	–	✓	✓	✓
PropertyDomain	✗	✓	✗	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
PropertyRange	✗	✓	✗	✓	✗	✓	✓	✓	✓	✓	✓	✓	✓	✓
SubProperty	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓
InverseProperty	–	–	–	–	–	✗	✗	–	–	–	✓	✓	✓	✓
EquivalentProperty	✗	✗	–	–	–	–	–	✓	–	–	–	✓	✓	✓
PropertyCharacteristic	✗	✗	✗	✗	✗	✗	✗	✓	✓	✓	✓	✓	✓	✓
PropertyChain	–	–	–	–	–	–	–	–	–	–	–	✓	✓	–

Lesson Learnt from Proof of Concept

Goal 1: showing JDSet datasets can be used by purely data-driven KGE models

Dataset	Best Model	*MRR*	*H@1*	*H@5*	*H@10*
DB100K-C	RotatE	0.371	0.282	0.468	0.541
YAGO3-39K-C	RotatE	0.225	0.140	0.298	0.391
YAGO4-20-C	RotatE	0.365	0.297	0.438	0.494
WHOW-5	RotatE	0.529	0.417	0.673	0.766
ARCO-10	RotatE	0.759	0.680	0.845	0.867
ERA-95	CompGCN	0.801	0.714	0.909	0.960

Goal 2: allow additional schema-driven validation (not possible with SOTA datasets))

- for plausible assertion (according to a score computed by a KGE model) run a consistency check (test executed with TransE on ARCO-10)
- **Results:** high model confidence do not necessarily imply semantic validity
  - the KG remains consistent for the first 1425 added predictions

Explanaining Link Predictions on KGs

by Tackling KG Semantics

Numeric-based methods consist of series of numbers without any obvious human interpretation

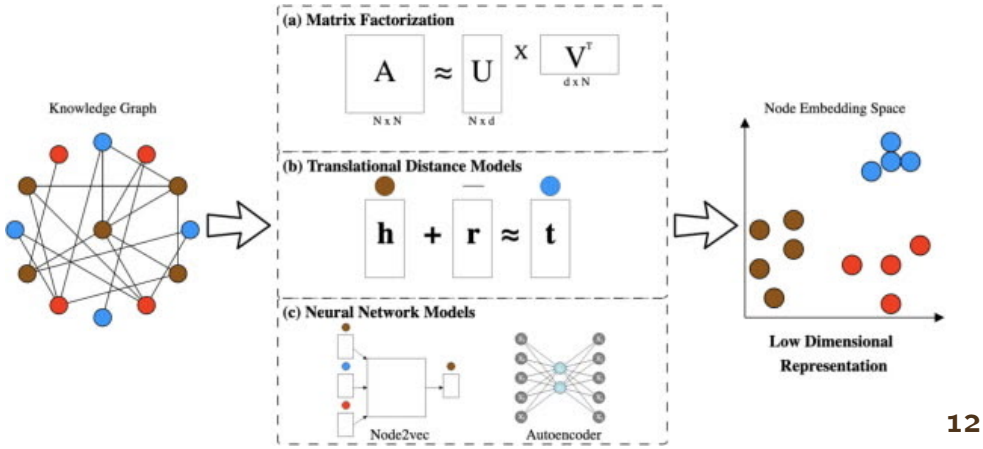

12

This may affects:

- the interpretability of the models
- the explainability of the results
- and possibly the trustworthiness of results

DRKG – Drug Repurposing Knowledge Graph

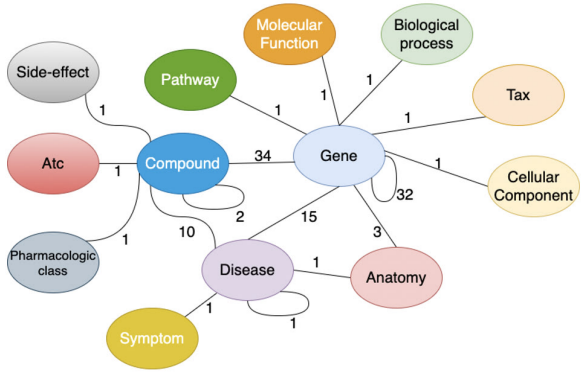

13

[12] 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

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

Computing Explanations: Approaches

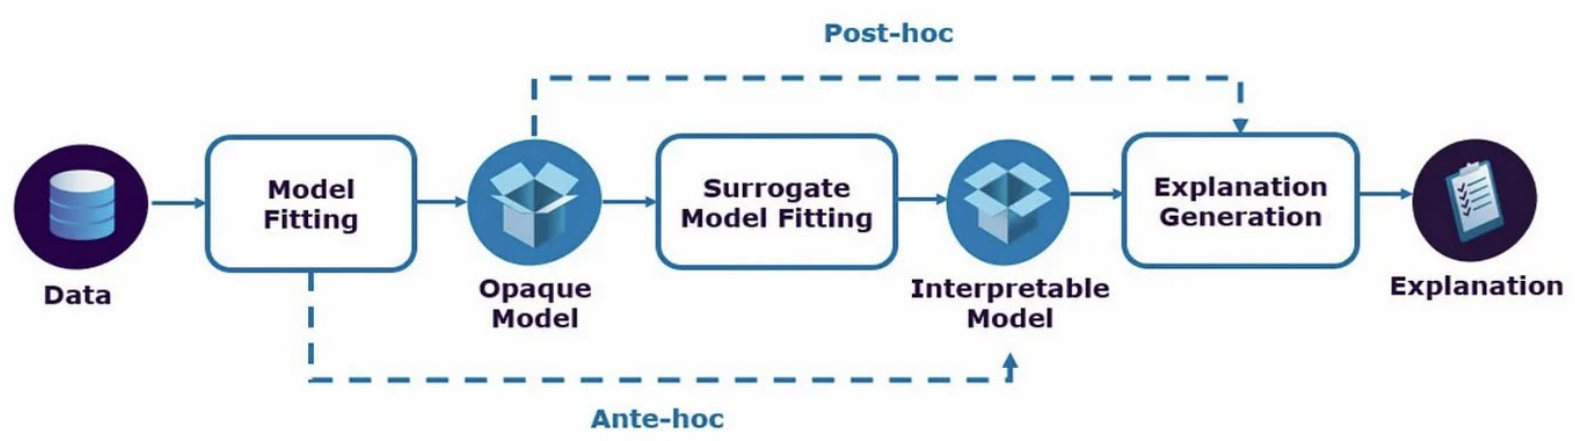

Explainable Machine Learning Approaches

14

Ante-hoc approach

- explain the model itself
- model-dependent

Post-hoc approach

- explain the result providing evidence from the data
- model-agnostic

[14] Picture from https://medium.com/@sparsha.stars/explainable-artificial-intelligence-technical-perspective-part-1-6eec91b4cc60

Post-hoc Explanations of Link Prediction (LP)

Post-hoc explanation methods are model agnostic. They find explanation(s) based on the output and the model input, independently on the KGE adopted

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

why is it provided?

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

Example of exmplanation

$\langle \texttt{NickMason}, \texttt{associatedBand}, \texttt{PinkFloyd} \rangle$,
$\langle \texttt{PinkFloyd}, \texttt{recordLabel}, \texttt{CapitolRecords} \rangle$

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

$\langle \texttt{RingoStarr}, \texttt{recordLabel}, \texttt{Parlophone} \rangle$

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

Post-hoc Explanations of LP: Current 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

- search for a path linking the subject s and object o of a predicted triple $\langle s, r, o \rangle$
  - **Max lenght 2** $\rightarrow$ six types of paths possible:
    Length 1: $P_1 = \{\langle s, r_q, o \rangle\}$, $P_2 = \{\langle o, r_q, s \rangle\}$
    Length 2: $P_3 = \{\langle e', r_q, s \rangle, \langle e', r', o \rangle\}$, $P_4 = \{\langle e', r_q, s \rangle, \langle o, r', e' \rangle\}$,
    $P_5 = \{\langle s, r_q, e' \rangle, \langle e', r', o \rangle\}$, $P_6 = \{\langle s, r_q, e' \rangle, \langle o, r', e' \rangle\}$,
    where r_q similar to r, r' any other relationship, e' any other entity;
- search driven by similarities between relation/entity embeddings via Euclidean distance
- structural comparisons with other paths in the KG to reinforce the reliability of the explanation found (referred to as support)

Exploiting Semantics for Providing Explanations to Link Predictions on KGs

KELPIE: High-Level Architecture...

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

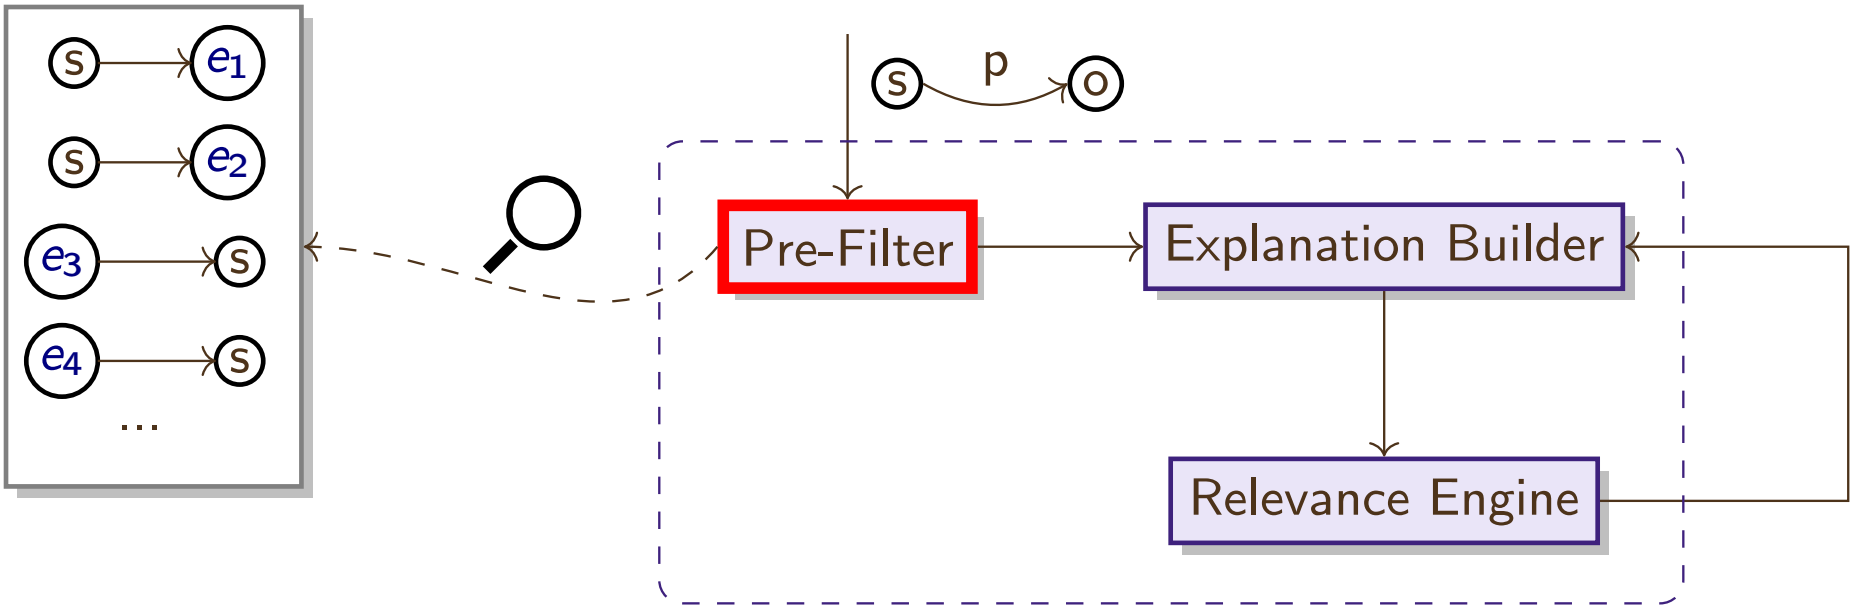

Pre-Filter:

- extracts the triples featuring s (either as subject or object)
- selects the k ones most topologically related to the prediction - structural similarity
- ignores semantic relations between facts

...KELPIE: High-Level Architecture

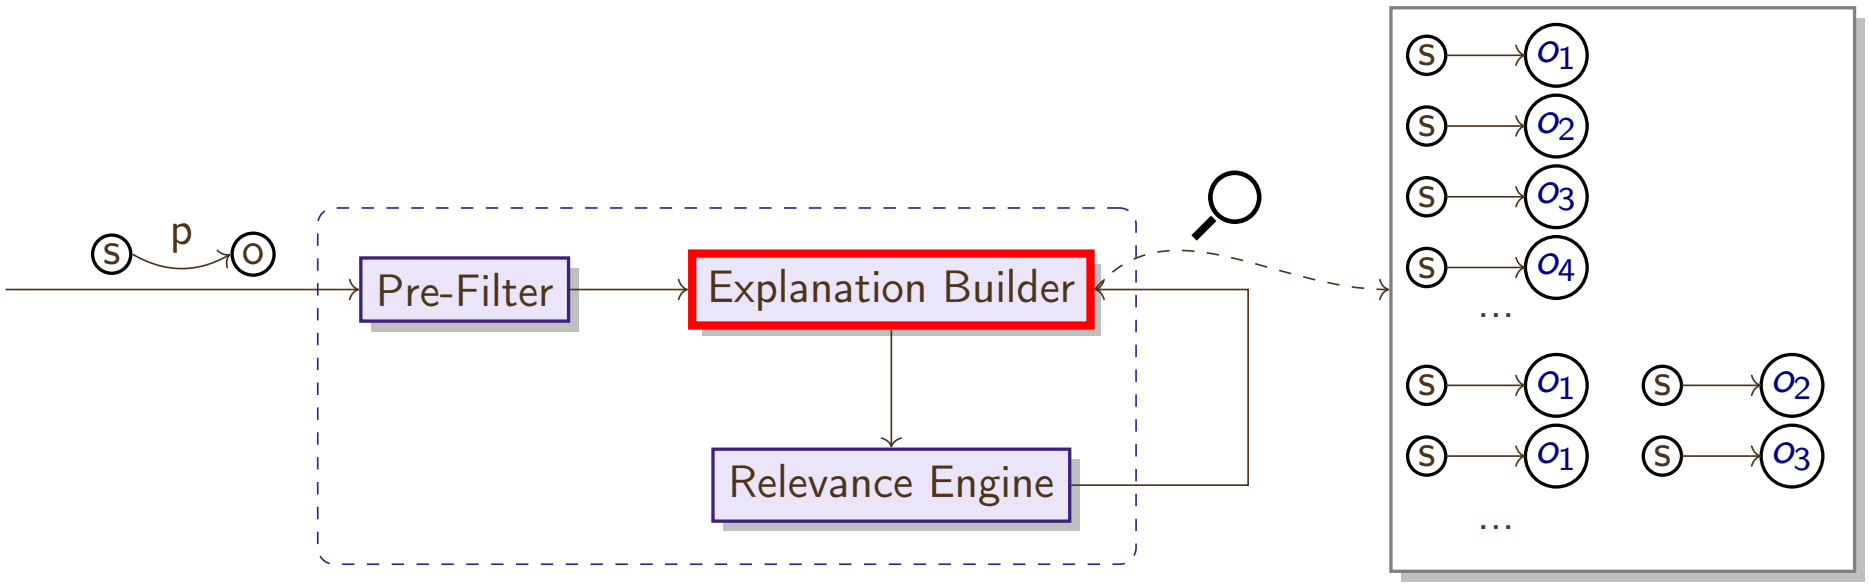

Explanation-Builder:

- combines the pre-filtered triples into candidate explanations
- leads the number of candidates to **explode combinatorially**

Relevance Engine: computes the relevance of candidate explanations

KELPIE++: Main features...

[Barile *et al.*2024b] [15]

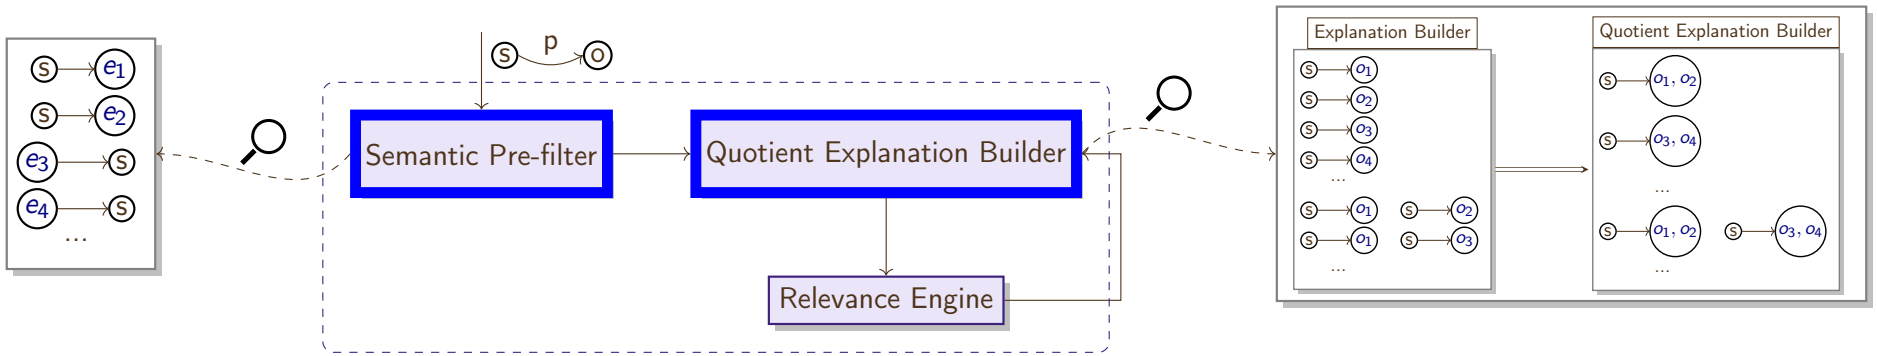

- Introduce Semantic Pre-filter by computing the Semantic Score
  - **Goal:** Improving the effectiveness/appropriateness of pre-filtered triples
- Introduce Quotient Explanation Builder of the pre-filtered triples
  - **Goal:** Improving the efficiency of the Explanation Builder

[15] R. Barile, C. d'Amato, N. Fanizzi. Explanation of link predictions on knowledge graphs via levelwise filtering and graph summarization. The Semantic Web - 21st International Conference, (ESWC 2024) Proceedings, Part I, volume 14664, pp. 26-30, LNCS Springer (2024)

...KELPIE++: Main features

[Barile *et al.*2024b] [16]

- Introduce Quotient Explanation Builder of the pre-filtered triples

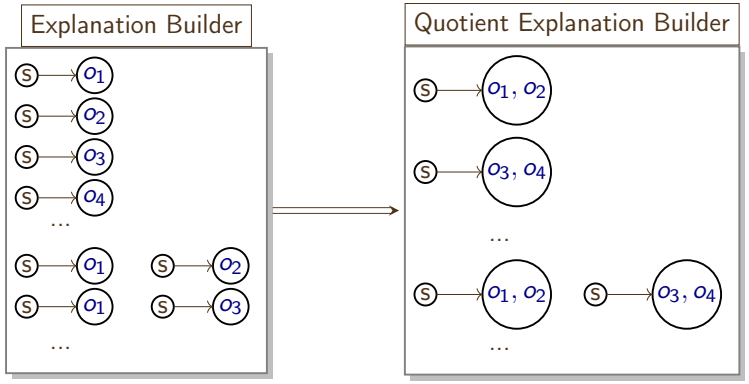



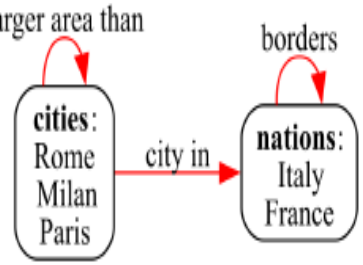

  - three different quotient graphs formulated
    - Simulation, Bisimulation, Depth-1 Bisimulation
    - exploit semantic information on types
    - offer different levels of granularity
  - combines quotient triples into candidate explanations
- Provides explanations at different level of details

[16] R. Barile, C. d'Amato, N. Fanizzi. Explanation of link predictions on knowledge graphs via levelwise filtering and graph summarization. The Semantic Web - 21st International Conference, (ESWC 2024) Proceedings, Part I, volume 14664, pp. 26-30, LNCS Springer (2024)

KELPIE++: explanations at different level of details

prediction: ⟨*Gabourey Sidibe*, *nationality*, *United States*⟩

dataset: YAGO4-20, **model**: ComplEx

KELPIE++ - Simulation quotient	
quotient triples	**entities**
⟨*Movie*, *actor*, *Gabourey Sidibe*⟩	*Casting By*, *Top Five*, *Seven Psychopaths*, *White Bird in a Blizzard*, *Tower Heist*, *Precious*
⟨*TVSeries*, *actor*, *Gabourey Sidibe*⟩	*Empire (2015 TV Series)*
KELPIE	
⟨*Empire (2015 TV Series)*, *actor*, *Gabourey Sidibe*⟩ ⟨*Top Five*, *actor*, *Gabourey Sidibe*⟩ ⟨*White Bird in a Blizzard*, *actor*, *Gabourey Sidibe*⟩	

Lesson Learnt from Experiments

Goal: Establishing the **impact of the Semantic Pre-Filter and Quotient Explanation Builder components when computing explanations** of link prediction results

- Comparing KELPIE++[17] to KELPIE as baseline
- TransE, ConvE, ComplEx adopted for the preliminary link prediction phase

KG	#Triples	#Entities	#Relationships
DBpedia50K	34289	24620	351
DBpedia100K	636802	98776	464
YAGO4-20	624580	96910	70

Results:

- KELPIE++ outperformed KELPIE in terms of
  - effectiveness (Δ HITS@1, Δ MRR)
  - efficiency (number of calls to the relevance engine)
  - particularly with **simulation** quotient graph

[17] Code and datasets publicly available https://github.com/rbarile17/kelpiePP

Evaluating Post-Hoc Explanations of Link Prediction on KGs

Evaluating Explanations: State of the Art and Issues

Approach 1: compare the computed explanations with ground truth

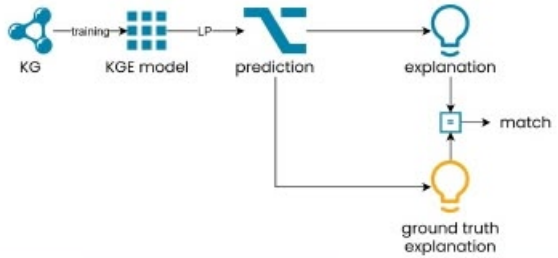

✓ • User aware
• Decoupled from the LP-X method

✗ Users recruiting

- Costly
  - needing recruiting users
- Hardly repeatable on new KG
  - ground truth missing

Approach 2: assess the impact-influence of explanations on solving the same LP task

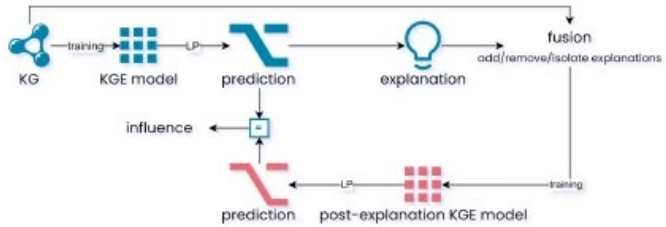

✓ Algorithmic

✗ • Not user aware
• Coupled with the LP-X method

- Different protocols adopted
  - Recall/Support vs. ΔHITS@1, ΔMRR
- Hard to compare different explanation solutions

Evaluation Explanation Protocol: Requirements

- **User-aware:** users can assess the utility of explanations
- **Algorithmic:** the evaluation of explanations can be automatized
  - user studies are costly
- **General:** explanations coming from different solutions can be compared

LP-DIXIT [Barile *et al.*2025] [18] ...

LP-DIXIT - the only protocol for evaluating explanations that is:

- user-aware
- fully algorithmic
- decoupled from the explanation method
- allows comparative analysis of different explanation methods - previously not possible

[18] R. Barile, C. d'Amato, N. Fanizzi. LP-DIXIT: Evaluating Explanations for Link Predictions on Knowledge Graphs using Large Language Models. Proceedings of the ACM on Web Conference 2025, (WWW 2025), pp. 4034–4042, ACM (2025)

...LP-DIXIT... [Barile *et al.*2025] [19]

LP-DIXIT grounded on forward simulatability cognitive theory [Hoffman *et al.*2023]

- a prediction is understandable if it is simulatable
- a prediction is simulatable if a verifier can guess its output given the same input
  - Forward Simulatability Variation (FSV) computed

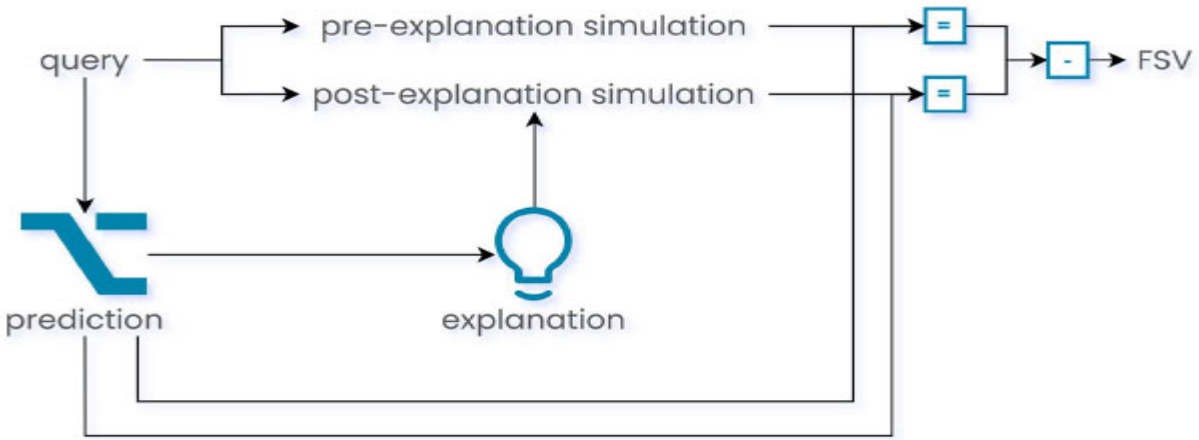

Large Language Models (LLMs) adopted for mimic the user(s) as verifier(s)

[19] R. Barile, C. d'Amato, N. Fanizzi. LP-DIXIT: Evaluating Explanations for Link Predictions on Knowledge Graphs using Large Language Models. Proceedings of the ACM on Web Conference 2025, (WWW 2025), pp. 4034–4042, ACM (2025)

...LP-DIXIT [Barile *et al.*2025] [21]

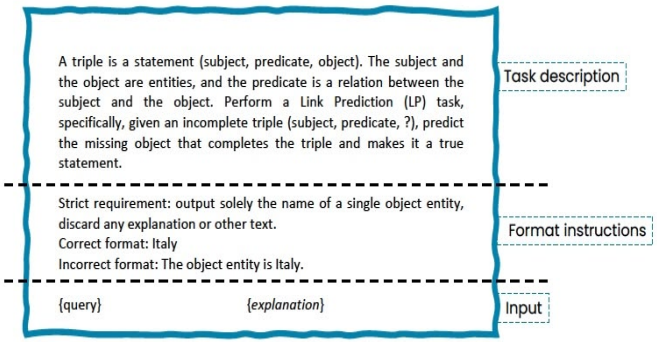

LP-DIXIT$_Z$: zero-shot without output constraint

task description	format instructions	input

LP-DIXIT$_O$: zero-shot with output constraint

task description	format instructions	candidate entities	input

- avoids that the LLMs output is not an entity in the KG.
- *candidate entities*: top-k ranked according to the LP method.

LP-DIXIT$_F$: few-shot without output constraint

task description	format instructions	demonstrations	input

- LLMs are not trained for LP on KGs
- demonstrations: triples with the same predicate of Q

LP-DIXIT$_F$: few-shot without output constraint

task description	format instructions	candidate entities	demonstrations	input

LP-DIXIT validated by:

- assessing its agreement with ground truth user judgements
- performing comparative evaluation of different explanation methods
  - previously not possible

GRainsaCK delivered: an open source library for benchmarking explanations [20]

[20] Code, datasets and documentation publicly available at https://github.com/rbarile17/grainsack

[21] R. Barile, C. d'Amato, N. Fanizzi. LP-DIXIT: Evaluating Explanations for Link Predictions on Knowledge Graphs using Large Language Models. Proceedings of the ACM on Web Conference 2025, (WWW 2025), pp. 4034–4042, ACM (2025)

Conclusions

Conclusions:

- NeSy LP and explanation solutions for KGs can be formalized taking into account semantics and reasoning capabilities
  - Framework for injecting semantics into KGE models
  - Solutions for injecting semantics when computing post-hoc explanations to LPs
- JediKG resource for systematic construction of datasets that combine schema and data from large and expressive KGs
- LP-DIXIT, a unified explanation evaluation protocol allowing the comparison of different methods

Next Research Challenges:

- Comprehensive evaluation of NeSy LP solutions on NeSy benchmarks
- Computation of semantic explanation via an integrated additive and subtractive couterfactual approach
- Extending the explanation evaluation protocol to multiple dimensions

Thank you



Roberto Barile



Ivan Diliso



Nicola Fanizzi



Nicola Flavio Quatraro

References

Barile, R., d'Amato, C., and Fanizzi, N. (2024a). Additive counterfactuals for explaining link predictions on knowledge graphs. In *he 24th International Conference on Knowledge Engineering and Knowledge Management, EKAW 2024, Amsterdam, Netherlands, November 26-28, 2024, Proceedings*, Lecture Notes in Computer Science. Springer.

Barile, R., d'Amato, C., and Fanizzi, N. (2024b). Explanation of link predictions on knowledge graphs via levelwise filtering and graph summarization. In A. Meroño-Peñuela, A. Dimou, R. Troncy, O. Hartig, M. Acosta, M. Alam, H. Paulheim, and P. Lisena, editors, *The Semantic Web - 21st International Conference, ESWC 2024, Hersonissos, Crete, Greece, May 26-30, 2024, Proceedings, Part I*, volume 14664 of *Lecture Notes in Computer Science*, pages 180–198. Springer.

Barile, R., d'Amato, C., and Fanizzi, N. (2025). LP-DIXIT: evaluating explanations for link predictions on knowledge graphs using large language models. In G. Long, M. Blumestein, Y. Chang, L. Lewin-Eytan, Z. H. Huang, and E. Yom-Tov, editors, *Proceedings of the ACM on Web Conference 2025, WWW 2025, Sydney, NSW, Australia, 28 April 2025- 2 May 2025*, pages 4034–4042. ACM.

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.

References

d'Amato, C., Quatraro, N. F., and Fanizzi, N. (2021a). 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. (2021b). 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.

d'Amato, C., Mahon, L., Monnin, P., and Stamou, G. (2023). Machine learning and knowledge graphs: Existing gaps and future research challenges. *TGDK*, **1**(1), 8:1–8:35.

Diliso, I., Barile, R., Fanizzi, N., and d'Amato, C. (2026). Jedikg: Return of the schema building complete datasets for machine learning and reasoning on knowledge graphs. In *ISWC Proceedings*, Lecture Notes in Computer Science. 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.

References

Hoffman, R. R., Mueller, S. T., Klein, G., and Litman, J. (2023). Measures for explainable AI: explanation goodness, user satisfaction, mental models, curiosity, trust, and human-ai performance. *Frontiers Comput. Sci.*, **5**.

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.

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.

Lourenço, B. F., Morgan, H., Ozaki, A., Pavlovic, A., and Sallinger, E. (2026). Boxlite: A faithful knowledge base embedding based on convex optimization (extended abstract). In *Proceedings of the 23rd International Conference on Principles of Knowledge Representation and Reasoning KR*, LNCS. Springer.

Minervini, P., Demeester, T., Rocktäschel, T., and Riedel, S. (2017). 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.

References

Robert, N., Monnin, P., and Faron, C. (2025). Semantically enriched datasets for link prediction: Db100k+, NELL-995+ and YAGO3-10+ (short paper). In A. Agiollo, E. Bardhi, G. Ciatto, S. Dumancic, and G. Marra, editors, *Proceedings of the 1st International Workshop on Advanced Neuro-Symbolic Applications co-located with the 28th European Conference on Artificial Intelligence (ECAI 2025), Bologna, Italy, October 26, 2025*, volume 4125 of *CEUR Workshop Proceedings*, pages 25–27. CEUR-WS.org.

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.

Xiong, B., Potyka, N., Tran, T.-K., Nayyeri, M., and Staab, S. (2022). Faithful embeddings for $\mathcal{EL}^{++}$ knowledge bases. In U. Sattler, A. Hogan, M. Keet, V. Presutti, J. P. A. Almeida, H. Takeda, P. Monnin, G. Pirrò, and C. d'Amato, editors, *The Semantic Web – ISWC 2022*, pages 22–38. Springer International Publishing.

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.