Towards Nonmonotonic Relational Learning from Knowledge Graphs

Hai Dang Tran[1], Daria Stepanova[1], Mohamed Gad Elrab[1], Francesca A. Lisi[2], Gerhard Weikum[1]

[1]Max Planck Institute for Informatics, Saarbrücken, Germany

[2]Università degli Studi di Bari "Aldo Moro", Bari, Italy

Motivation

- Knowledge Graphs: huge collections of $\langle$*subject predicate object*$\rangle$ triples

  $\langle$*bob isMarriedTo alice*$\rangle$, $\langle$*alice type researcher*$\rangle$

- Encode positive unary/binary facts under Open World Assumption (OWA)

  $isMarriedTo(bob, alice), researcher(alice)$

- KGs are automatically constructed, possibly incomplete and inaccurate

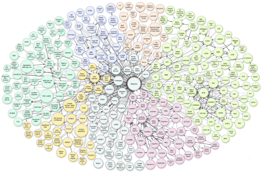

Motivation

Horn rule mining to complete KGs, [Galárraga *et al.*, 2015]

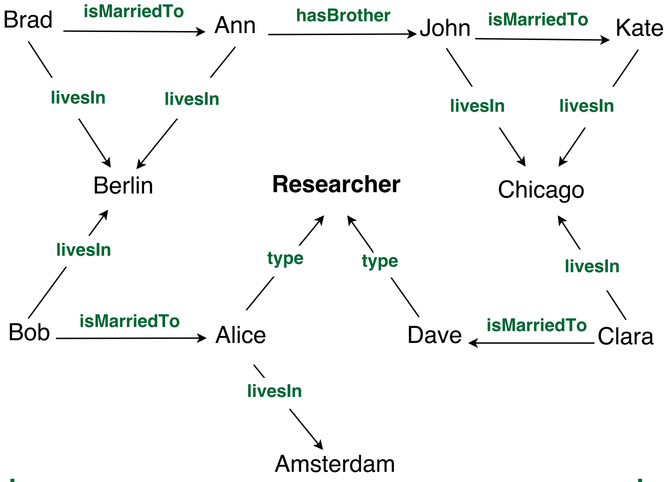

$$r: \; livesIn(X, Z) \leftarrow isMarriedTo(Y, X), livesIn(Y, Z)$$

Motivation

Horn rule mining to complete KGs, [Galárraga *et al.*, 2015]

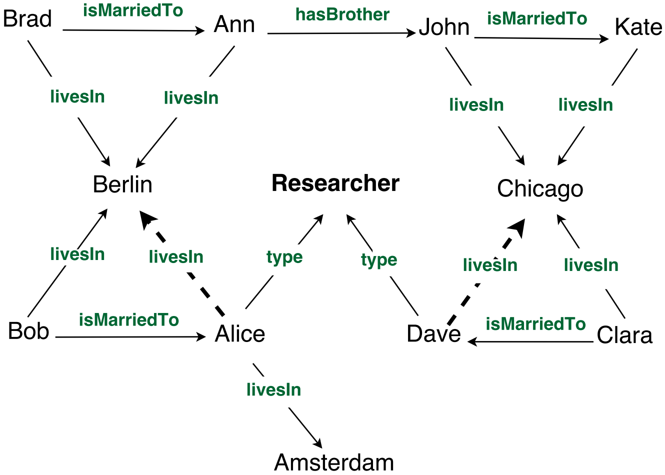

$$r: \ livesIn(X,Z) \leftarrow isMarriedTo(Y,X), livesIn(Y,Z)$$

Motivation

In this work: nonmonotonic rule learning on KGs, OWA is a challenge!

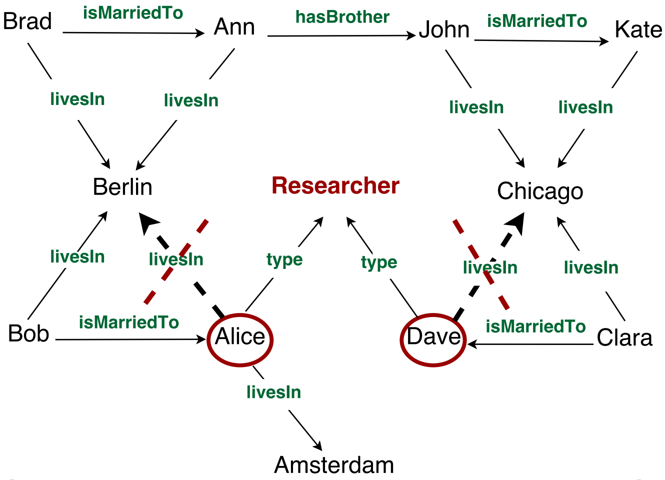

$$r: \ livesIn(X,Z) \leftarrow isMarriedTo(Y,X), livesIn(Y,Z), not \ \ researcher(X)$$

Problem Statement

ILP-based theory revision under CWA [Wrobel, 1996], . . .

Quality-based Horn Theory Revision (QHTR)

Given:

- KG $\mathcal{G}$
- Horn ruleset $\mathcal{R}_H$

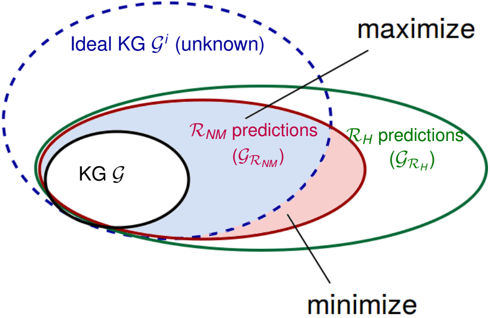

Find:

- nonmonotonic revision $\mathcal{R}_{NM}$ of $\mathcal{R}_H$, such that its predictive quality is better then of $\mathcal{R}_H$

Conflicting Predictions

Ensure quality of exceptions by minimizing conflicts

$$\mathcal{R}_{NM}^{aux} = \left\{ \begin{array}{l} r1:\ livesIn(X,Z) \leftarrow isMarTo(Y,X), livesIn(Y,Z), not\ \ res(X) \\ r1^{aux}:\ not\_livesIn(X,Z) \leftarrow isMarTo(Y,X), livesIn(Y,Z), res(X) \\ r2:\ livesIn(X,Z) \leftarrow bornIn(X,Z), not\ \ immigrant(X) \\ r2^{aux}:\ not\_livesIn(X,Z) \leftarrow bornIn(X,Z), immigrant(X) \end{array} \right\}$$

$\{livesIn(c,d), not\_livesIn(c,d)\} \in \mathcal{G}_{\mathcal{R}_{NM}^{aux}}$ are conflicting predictions

Intuition: researcher might be a strong exception for *r1*, but application of *r2* to the KG could weaken it; less conflicts less weak exceptions

Problem Statement

Quality-based Horn Theory Revision (QHTR)

Given:

- KG $\mathcal{G}$
- Horn ruleset $\mathcal{R}_H$

Find:

- nonmonotonic revision $\mathcal{R}_{NM}$ of $\mathcal{R}_H$, such that
  - number of conflicting predictions made by $\mathcal{R}_{NM}^{aux}$ is **minimal**
  - average conviction $conv(r, \mathcal{G}) = \dfrac{1 - supp(r, \mathcal{G})}{1 - conf(r, \mathcal{G})}$ is **maximal** [Azevedo and Jorge, 2007]

Related Work

- **First-order theory revision**
  - RUTH [Adé *et al.*, 1994]
  - FORTE [Richards and Mooney, 1995]

  . . .

- **Learning nonmonotonic programs**
  - [Dimopoulos and Kakas, 1995]
  - ILASP [Law *et al.*, 2015]
  - ILED [Katzouris *et al.*, 2015]

  . . .

- **Outlier detection in logic programs**
  - [Angiulli and Fassetti, 2014]

  . . .

- **Mining rules with exceptions**
  - [Suzuki, 2006]

  . . .

Approach Overview

Extension of our results from [Gad-Elrab *et al.*, 2016] to binaries

Step 1. Mine predictive association rules in the form of first-order Horn clauses, [Galárraga *et al.*, 2015]

Step 2. Determine normal and abnormal substitutions for every $r \in \mathcal{R}_H$

Step 3. Find all exception candidates for every rule

Step 4. Rank exception candidates and select the locally best ones

Step 2: (Ab)normal Substitutions

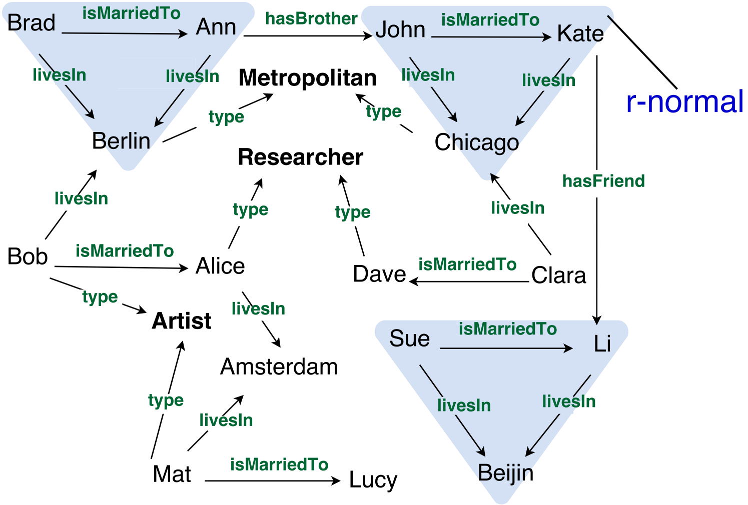

$$r: \ livesIn(X, Z) \leftarrow isMarriedTo(Y, X), livesIn(Y, Z)$$

Step 2: (Ab)normal Substitutions

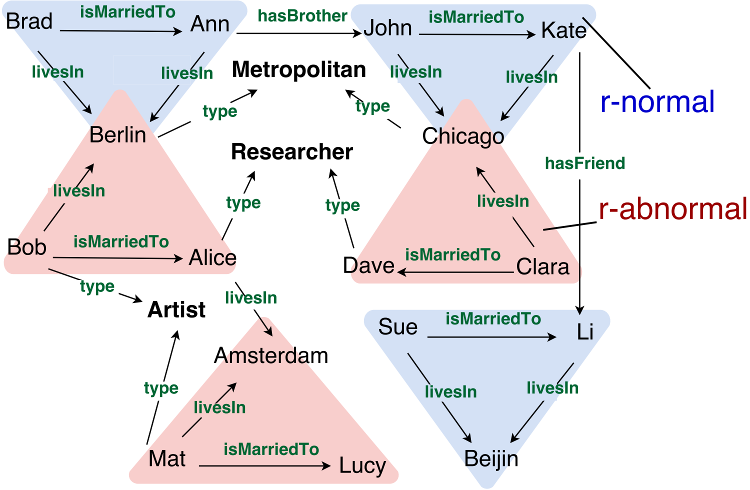

$$r: \; livesIn(X, Z) \leftarrow isMarriedTo(Y, X), livesIn(Y, Z)$$

Step 3: Exception Candidates

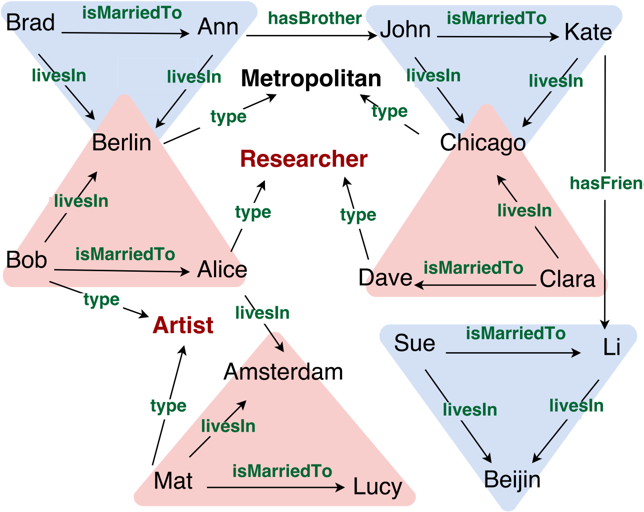

$r: \ livesIn(X,Z) \leftarrow isMarriedTo(Y,X), livesIn(Y,Z)$ $not \ researcher(X)$ $not \ artist(Y)$

Step 4: Exception Ranking

$$r1 \ldots\ldots\ldots \{\underline{\mathbf{e_1}}, e_2, e_3, \ldots\}$$
$$r2 \ldots\ldots\ldots \{e_1, \underline{\mathbf{e_2}}, e_3, \ldots\}$$
$$r3 \ldots\ldots\ldots \{\underline{\mathbf{e_1}}, e_2, e_3, \ldots\}$$

Finding globally best revision is expensive, too many candidates!

- Naive ranking: pick for $r \in \mathcal{R}_H$ a revision r' with the highest $conv(r, \mathcal{G})$
- Partial materialization: first materialize all rules apart from r, then pick a revision with the highest $\dfrac{conv(r, \mathcal{G}') + conv(r^{aux}, \mathcal{G}')}{2}$
- Ordered part. mat. (OPM): same as part. mat., but materialize only rules ordered higher then r based on *conv*

Preliminary Experiments

- $\mathcal{G}^i_{appr}$: IMDB (movie) KG[1]: $\approx$ 600.000 facts, $\approx$ 40 relations
  E.g., *directedBy*, *actedIn*
- $\mathcal{G}$: random. rem. 20% from $\mathcal{G}^i_{appr}$ for every relation
- $\mathcal{R}_H$: $h(X, Y) \leftarrow p(X, Z), q(Z, Y)$ mine from $\mathcal{G}$
- Exception types: $e_1(X), e_2(Y), e_3(X, Y)$
- OPM ranker, predictions are computed by answer set solver dlv[2]

---

[1] `http://imdb.com`
[2] `http://dlvsystem.com`

Preliminary Experiments

k	avg. conv.		confl.	number of predictions					
				$\mathcal{R}_H$		$\mathcal{R}_{NM}$		$\mathcal{R}_H$ **not** $\mathcal{R}_{NM}$	
	$\mathcal{R}_H$	$\mathcal{R}_{NM}$	$\mathcal{R}_{NM}$	**all**	**in** $\mathcal{G}^i_{appr}$	**all**	**in** $\mathcal{G}^i_{appr}$	**false** ✓	**in** $\mathcal{G}^i_{appr}$
5	4.08	6.16	0.28	345	161	331	156	0	14
10	2.91	4.21	0.08	2178	456	2118	450	27	33
15	2.5	3.42	0.09	3482	629	3348	622	86	48
20	2.29	3.0	0.13	5278	848	5046	835	157	75

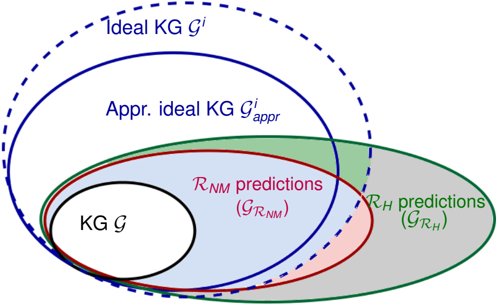

Preliminary Experiments

k	avg. conv.		confl.	number of predictions					
				$\mathcal{R}_H$		$\mathcal{R}_{NM}$		$\mathcal{R}_H$ **not** $\mathcal{R}_{NM}$	
	$\mathcal{R}_H$	$\mathcal{R}_{NM}$	$\mathcal{R}_{NM}$	**all**	**in** $\mathcal{G}^i_{appr}$	**all**	**in** $\mathcal{G}^i_{appr}$	**false** ✓	**in** $\mathcal{G}^i_{appr}$
5	4.08	6.16	0.28	345	161	331	156	0	14
10	2.91	4.21	0.08	2178	456	2118	450	27	33
15	2.5	3.42	0.09	3482	629	3348	622	86	48
20	2.29	3.0	0.13	5278	848	5046	835	157	75

Table : Top k rule revision results

Examples of revised rules:

r_1 : *writtenBy*$(X, Z) \leftarrow$ *hasPredecessor*(X, Y), *writtenBy*(Y, Z), **not** *is_American_film*(X)

r_2 : *actedIn*$(X, Z) \leftarrow$ *isMarriedTo*(X, Y), *directed*(Y, Z), **not** *is_silent_film_actor*(X)

Summary

Contributions:

- Quality-based Horn theory revision framework under OWA
- Approach for computing and ranking exceptions based on partial materialization
- Preliminary experiments on a real-world KG

Further Work:

- Evidence for and against exceptions from text corpora
- Partial completeness
- Causality of rules, probabilities
- More complex rules, e.g. with existentials

References I

Hilde Adé, Bart Malfait, and Luc De Raedt.
RUTH: an ILP theory revision system.
In *Methodologies for Intelligent Systems, 8th International Symposium, ISMIS '94, Charlotte, North Carolina, USA, October 16-19, 1994, Proceedings*, pages 336–345, 1994.

Fabrizio Angiulli and Fabio Fassetti.
Exploiting domain knowledge to detect outliers.
Data Min. Knowl. Discov., 28(2):519–568, 2014.

Paulo J. Azevedo and Alípio Mário Jorge.
Comparing Rule Measures for Predictive Association Rules.
In *Proceedings of ECML*, pages 510–517, 2007.

Yannis Dimopoulos and Antonis C. Kakas.
Learning non-monotonic logic programs: Learning exceptions.
In *Machine Learning: ECML-95, 8th European Conference on Machine Learning, Heraclion, Crete, Greece, April 25-27, 1995, Proceedings*, pages 122–137, 1995.

M. Gad-Elrab, D. Stepanova, J. Urbani, and G. Weikum.
Exception-enriched Rule Learning from Knowledge Graphs.
In *In Proc. of ISWC 2016, to appear*, 2016.

Luis Galárraga, Christina Teflioudi, Katja Hose, and Fabian M. Suchanek.
Fast Rule Mining in Ontological Knowledge Bases with AMIE+.
In *VLDB Journal*, 2015.

Nikos Katzouris, Alexander Artikis, and Georgios Paliouras.
Incremental learning of event definitions with inductive logic programming.
Machine Learning, 100(2-3):555–585, 2015.

Mark Law, Alessandra Russo, and Krysia Broda.
The ILASP system for learning answer set programs, 2015.

References II

Bradley L. Richards and Raymond J. Mooney.
Automated refinement of first-order horn-clause domain theories.
Machine Learning, 19(2):95–131, 1995.

Einoshin Suzuki.
Data mining methods for discovering interesting exceptions from an unsupervised table.
J. UCS, 12(6):627–653, 2006.

Stefan Wrobel.
First order theory refinement.
In Luc De Raedt, editor, *Advances in Inductive Logic Programming*, pages 14 — 33. IOS Press, Amsterdam, 1996.