Open Relation Modeling: Learning to Define Relations between Entities

Jie Huang, Kevin Chen-Chuan Chang, Jinjun Xiong, Wen-mei Hwu

University of Illinois at Urbana-Champaign
University at Buffalo
NVIDIA

NVIDIA

Relationships exist widely

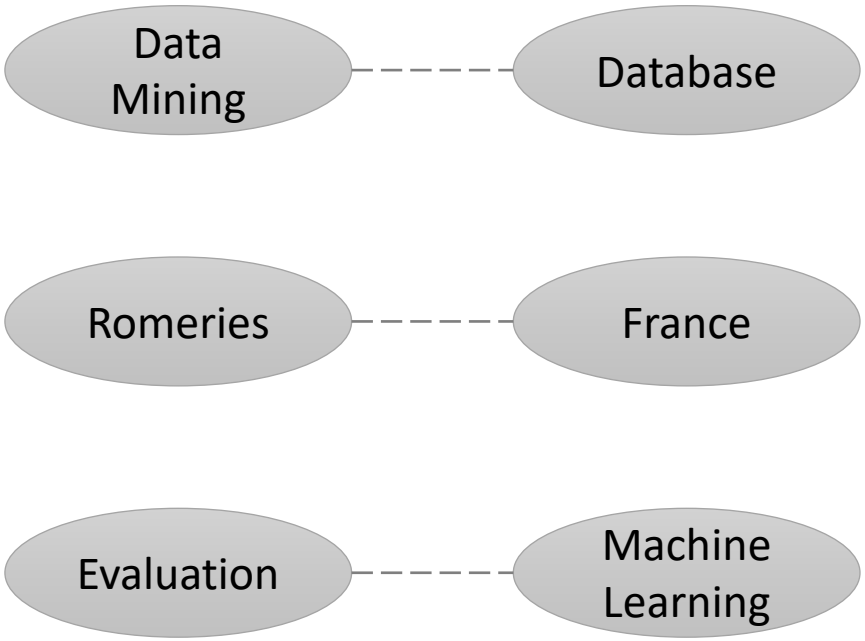

To represent relationships…

A fact/reasoning path in KG: (data mining, *facet-of*, database)
A sentence: "we study data mining and database."

***Not interpretable*:** cannot know exactly how they are related
Not open: may not exist a fact or a sentence containing them

Open Relation Modeling

Open Relation Modeling: given two entities, generating a coherent sentence describing the relationship between them, where types of relations do not need to be pre-specified.

E.g., "***data mining*** is a process of extracting and discovering patterns in large data sets involving methods at the intersection of machine learning, statistics, and ***database*** systems."

Interpretable & Open!

Open Relation Modeling: Learning from definitions

"***data mining*** is a process of extracting and discovering patterns in large data sets involving methods at the intersection of machine learning, statistics, and ***database*** systems."

Definitions of entities: informative sentences that capture the most representative characteristics of entities

Find the relation between two entities by defining one entity in terms of the other entity!

Open Relation Modeling: Learning from definitions

- Extract entity pairs from definitions of entities
- Fine-tune BART (Lewis et al., 2020a) to reproduce the definitions of entities with extracted entity pairs as input

Haste is a municipality in the district of Schaumburg, in Lower Saxony, Germany. → (Haste, Schaumburg) (Haste, Germany) …

(Haste, Germany) → **BART** → ***Haste*** is a municipality in the district of Schaumburg, in Lower Saxony, ***Germany***.

=> **RelationBART-Vanilla**

Reasoning Path-Enriched Relation Modeling

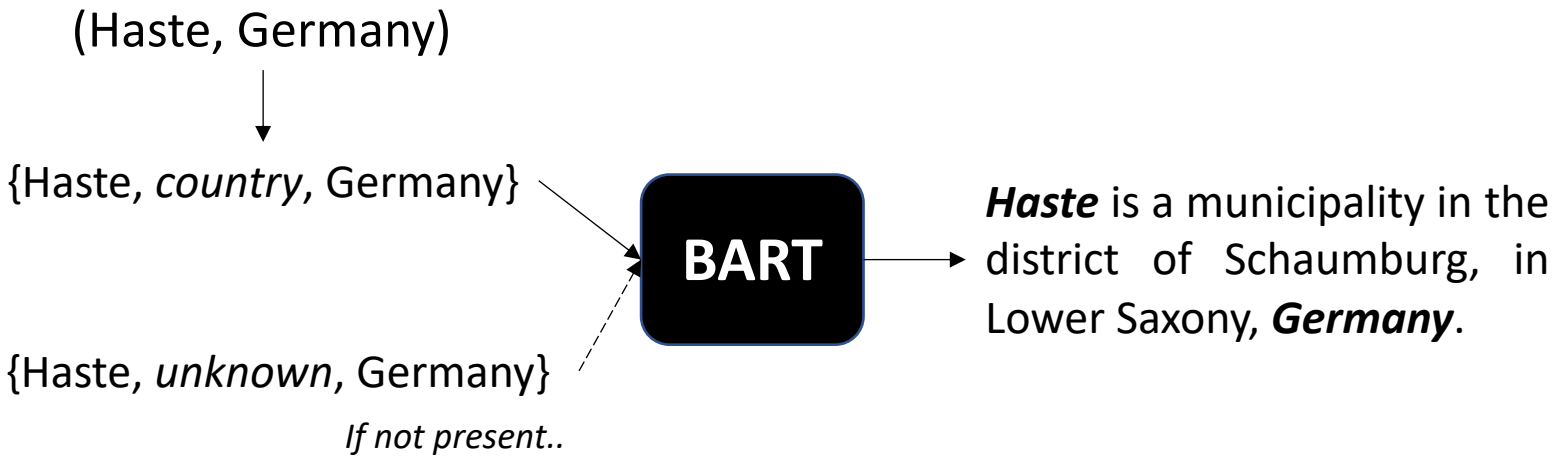

Augment the input with the shortest reasoning path => **RelationBART-SP**

Open Relation Modeling with Reasoning Path Selection

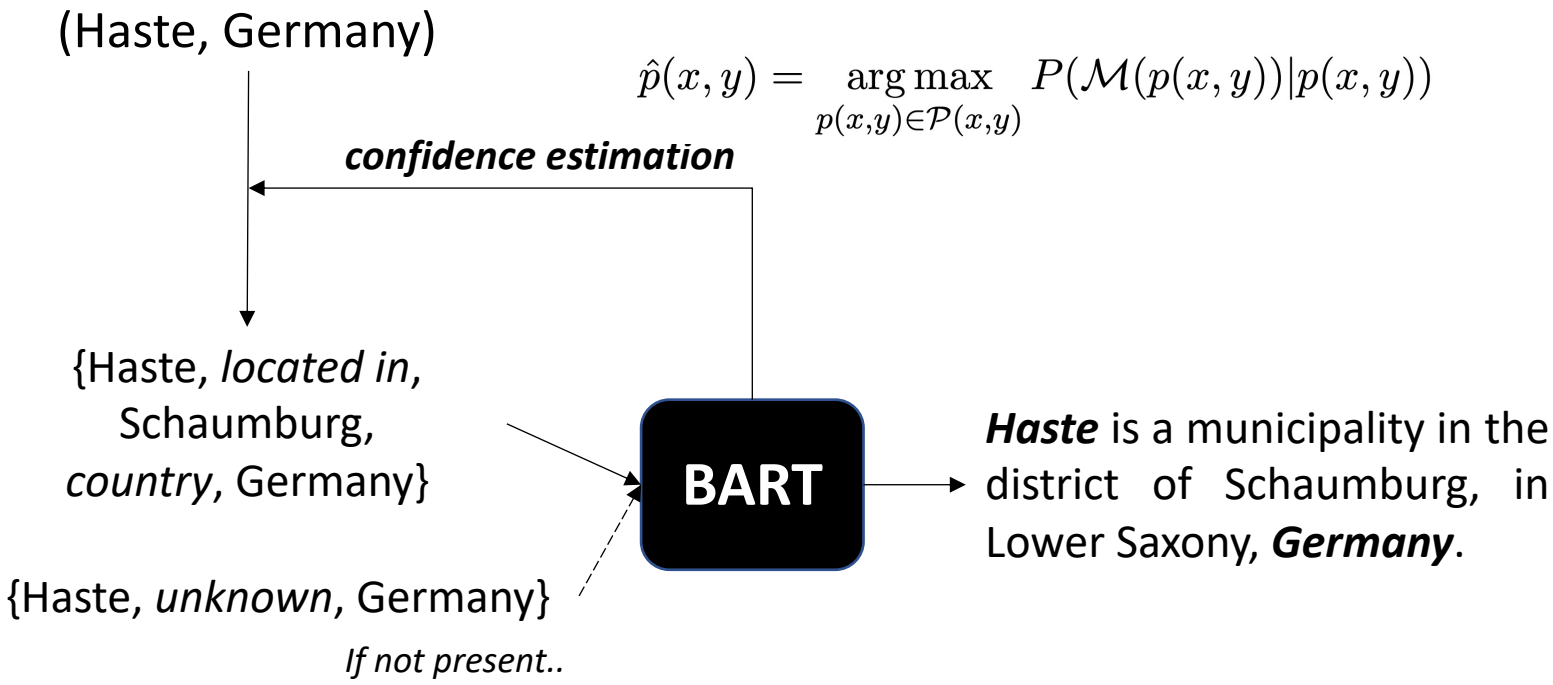

Select the best reasoning path with confidence estimation
=> **RelationBART-SP/MP + PS**

Experiments: Dataset

	train	dev	test	test*
number	5,434,158	27,431	55,226	7,302

	1-hop	2-hop	3-hop	> 3-hop
ratio (%)	35.14	17.80	7.33	39.73

Entity pairs -> First sentences of Wikipedia pages

*test** denotes a filtered sub-test set with a higher quality

Experiments: Results

Quantitative

	BL	R-L	MT	BS
DefBART	25.98	47.38	22.39	83.41
RelationBART-Vanilla (w/o PT)	34.70	59.57	28.85	88.01
RelationBART-SP (w/o PT)	35.48	60.55	29.40	88.43
RelationBART-SP (w/o PT) + PS	38.62	62.60	31.07	89.05
RelationBART-MP (w/o PT) + PS	**40.52**	**63.73**	**32.06**	**89.53**
RelationBART-Vanilla	35.45	59.92	29.33	88.25
RelationBART-SP	36.58	61.15	30.04	88.75
RelationBART-SP + PS	39.93	63.32	31.80	89.39
RelationBART-MP + PS	**41.43**	**64.15**	**32.45**	**89.64**
RelationBART-Vanilla (Large)	36.53	60.54	29.90	88.50
RelationBART-SP (Large)	37.65	61.34	30.57	88.89
RelationBART-SP (Large) + PS	41.21	63.56	32.41	89.53
RelationBART-MP (Large) + PS	**41.46**	**64.36**	**32.62**	**89.79**

Qualitative

	Rating (1-4)
RelationBART-Vanilla (Large)	2.67
RelationBART-SP (Large)	2.82
RelationBART-MP (Large) + PS	**3.01**

Refer to the paper for more results and details.

Experiments: Generation Examples

Input	Reasoning Path (Encoding)	Output
Romeries **[reference]**	-	Romeries is a commune in the Nord department in northern France.
(Romeries, France)	Romeries; country: France **[NOT selected]**	Romeries is a commune in the Somme department in Hauts-de-France in northern France.
	Romeries; located in the administrative territorial entity: Nord; country: France **[selected]**	Romeries is a commune in the Nord department in northern France.
	Romeries; shares border with: Solesmes; country: France **[NOT selected]**	Romeries is a commune in the Marne department in north-eastern France.
data mining **[reference]**	-	Data mining is a process of extracting and discovering patterns in large data sets involving methods at the intersection of machine learning, statistics, and database systems.
(data mining, artificial intelligence)	data mining; facet of: machine learning; subclass of: artificial intelligence	Data mining is a subfield of machine learning and artificial intelligence concerned with the collection, processing, and analysis of large amounts of data.
(data mining, database)	data mining; facet of: database	Data mining is the process of extracting information from a data set, such as a database.
(data mining, statistics)	data mining; facet of: machine learning; partially coincident with: statistics	Data mining is the process of collecting, processing, and analyzing large amounts of data in the form of statistics and machine learning.
(data mining, data)	data mining; facet of: machine learning; partially coincident with: statistics; studies: data	Data mining is the process of extracting useful information from data.
evaluation **[reference]**	-	Evaluation is a systematic determination of a subject's merit, worth and significance, using criteria governed by a set of standards.
(evaluation, machine learning)	evaluation; unknown: machine learning	In computer science, evaluation is the process of evaluating a machine learning algorithm to determine whether the algorithm is performing well.
(evaluation, algorithm)	evaluation; unknown: algorithm	In computer science, evaluation is the process of evaluating an algorithm to determine whether it is correct.
(evaluation, robotics)	evaluation; unknown: robotics	In robotics, evaluation is the process of determining whether or not a particular component of a system is working properly.
(evaluation, software engineering)	evaluation; unknown: software engineering	In computer science and software engineering, evaluation is the process of determining whether a particular feature or feature should be added to a product or service.
(evaluation, computer security)	evaluation; unknown: computer security	In computer security, evaluation is the process of determining the security of a computer system.

Conclusion

- We introduce ***open relation modeling*** – generating coherent sentences describing entity relationships;
- To solve this task, we propose to teach machines to generate definition-like relation descriptions by letting them learn from producing definitions conditioned on extracted entity pairs;
- We apply PLMs and design reasoning path-enriched PLMs for open relation modeling;
- Experimental results show that our methods can generate sentences that well capture entity relationships.

Email: jeffhj@illinois.edu
Code and data: https://github.com/jeffhj/open-relation-modeling

Thanks!