Learning Relation Entailment with Structured and Textual Information

Zhengbao Jiang[1], Jun Araki[2], Donghan Yu[1], Ruohong Zhang[1],
Wei Xu[3], Yiming Yang[1], Graham Neubig[1]

Carnegie Mellon University[1], Bosch Research North America[2], Ohio State University[3]

zhengbaj@cs.cmu.edu

Motivation

- Relations among entities play a fundamental role in knowledge graphs.

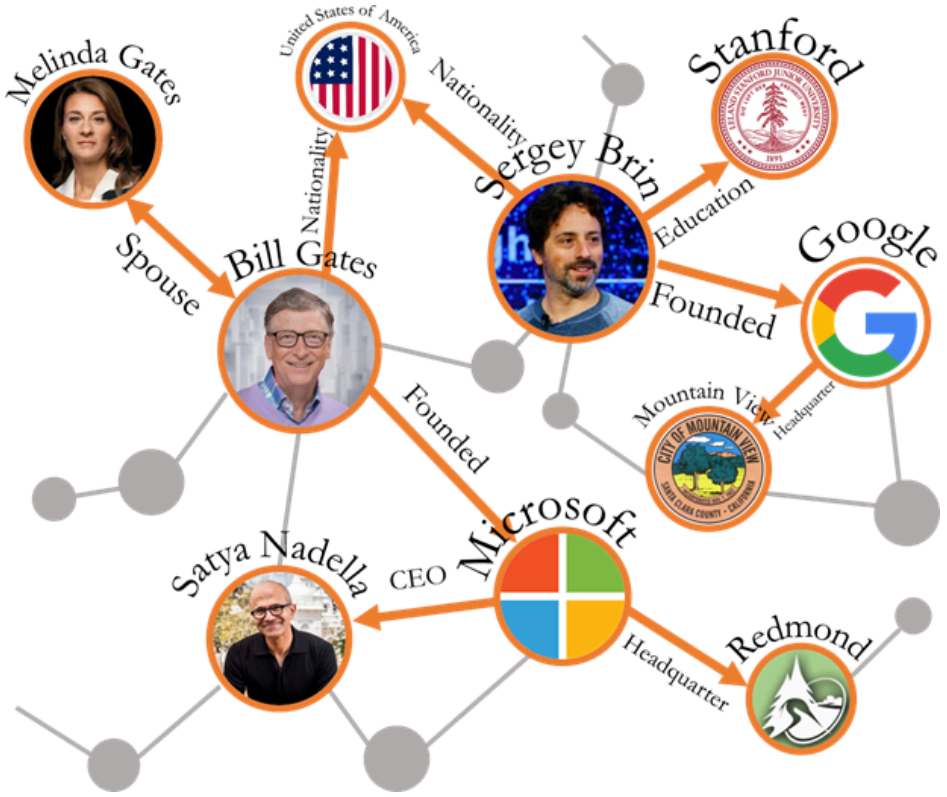

Motivation

- However, relations are treated as independent.

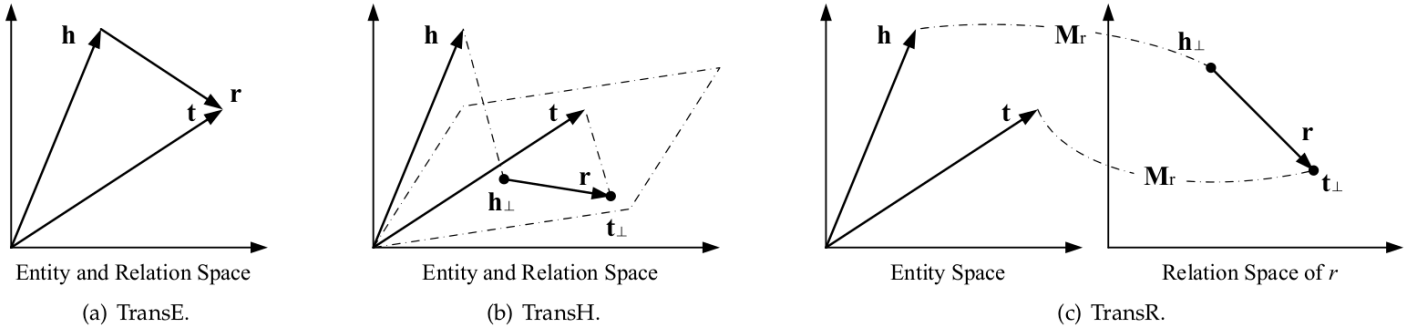

KG embedding: each relation is treated as an atomic unit with separate parameters.

Motivation

- However, relations are treated as independent.

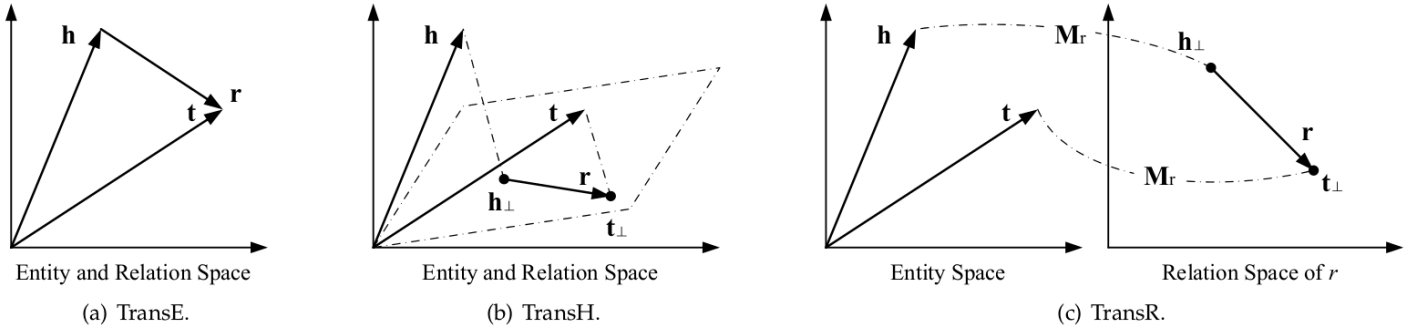

KG embedding: each relation is treated as an atomic unit with separate parameters.

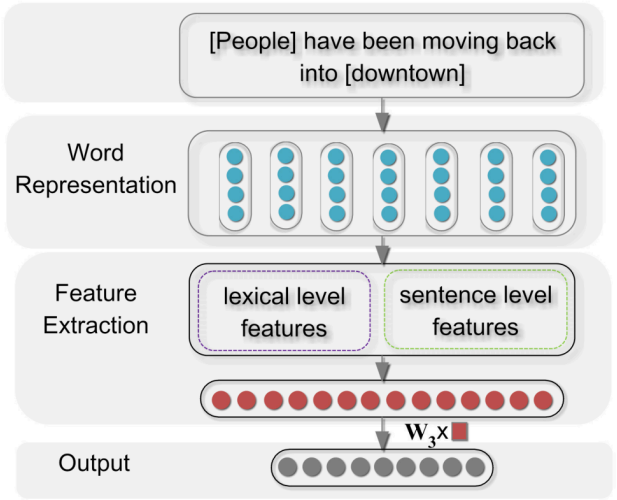

Relation extraction: each relation is an independent class.

Meta-relation: Relations Between Relations

- Relation entailment: existence of one relation can entail the existence of another relation.

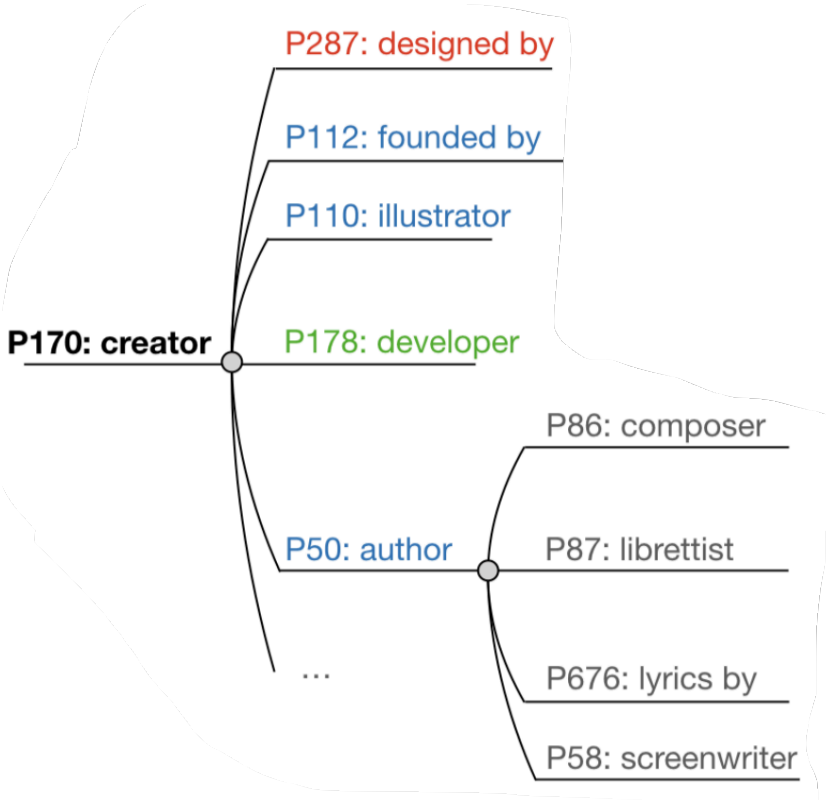

Applications of Relation Entailment

- Knowledge graph representation learning.

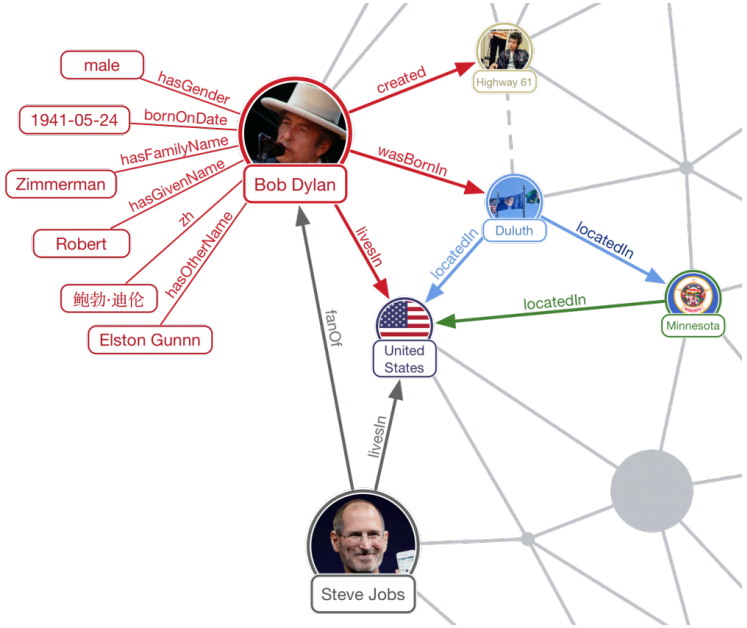

Applications of Relation Entailment

- Knowledge graph representation learning.
- Relation extraction.

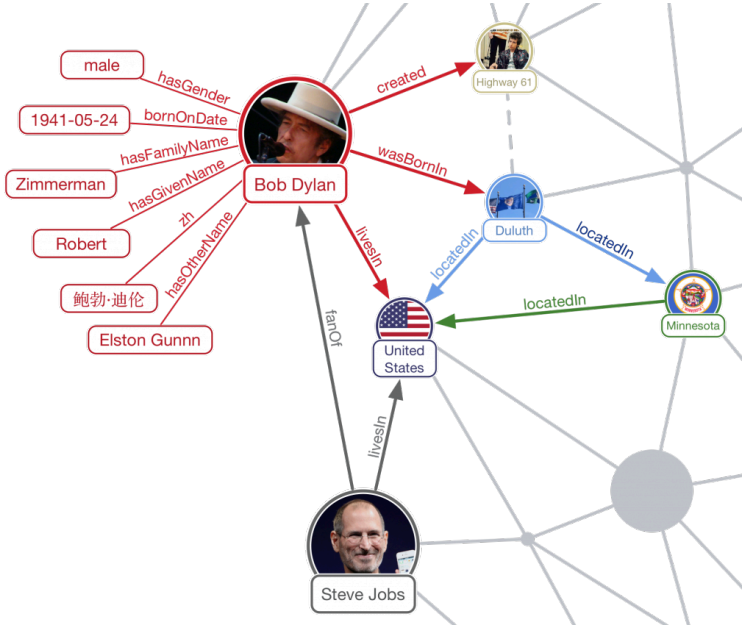

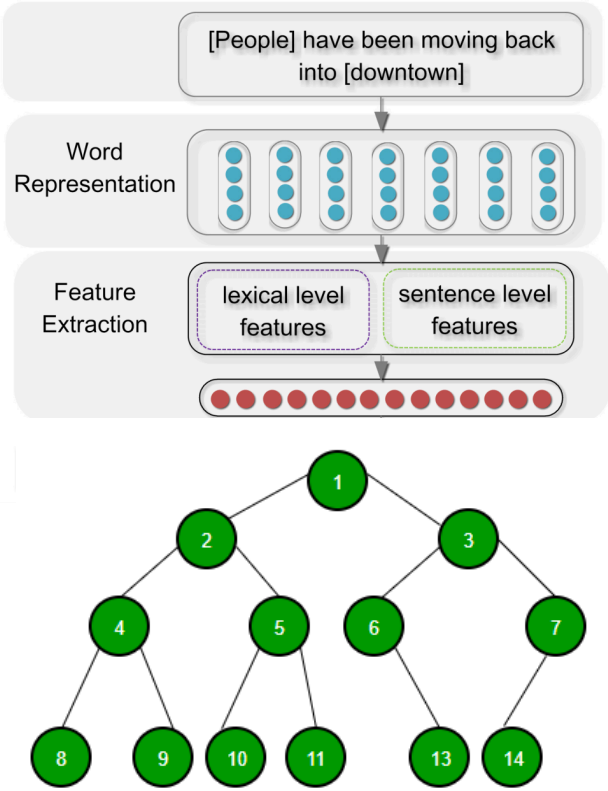

Applications of Relation Entailment

- Knowledge graph representation learning.
- Relation extraction.
- KG-based question answering.

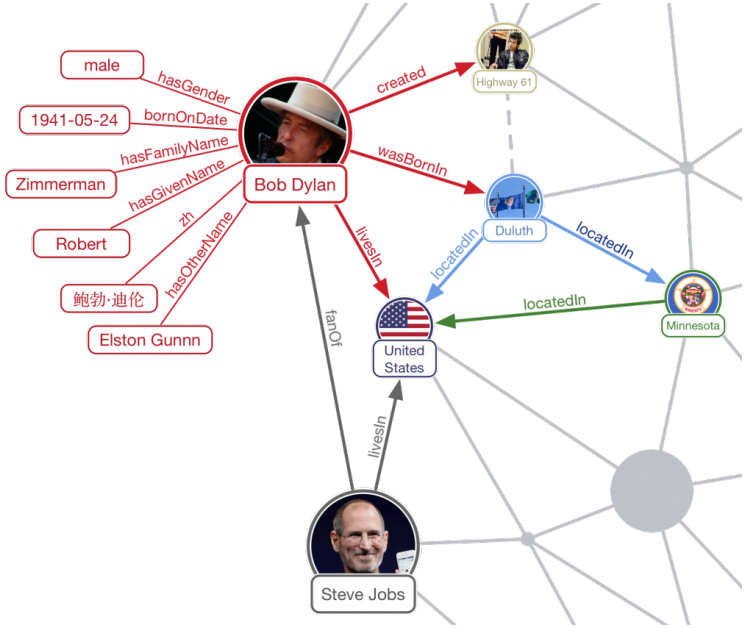

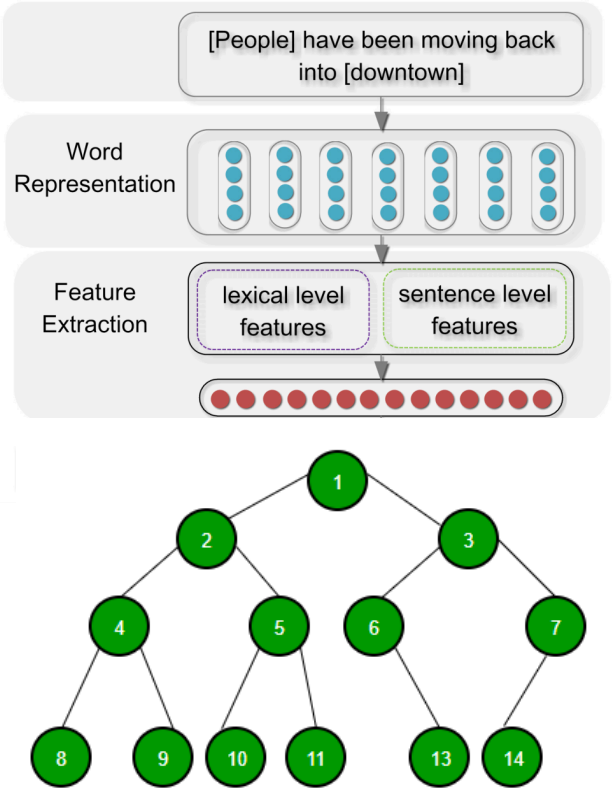

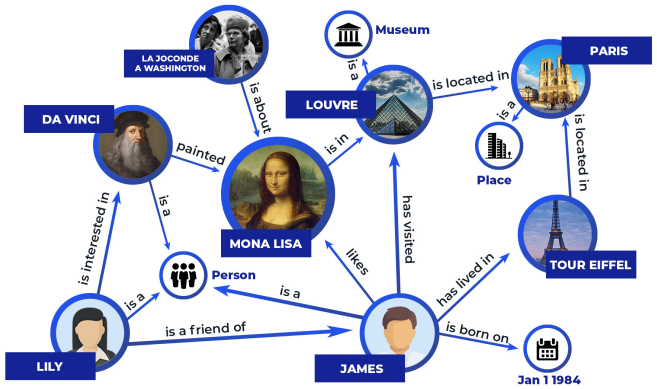

Relation Entailment Task Definition

h	t
(Apple,	Jobs)
(Microsoft,	Gates)
(Facebook,	Zuckerberg)
…	

r=founded by

Relation Entailment Task Definition

- Notations
  - Head and tail entities $h, t \in \mathcal{E}$.
  - Relations $r \in \mathcal{R}$.
  - Instances of a relation $C_r = \{(h, r, t)^{(i)}\}_i$.
- Relation entailment
  - $r \vDash r'$ if and only if $C_r \subseteq C_{r'}$.
- Task of predicting relation entailment
  - Given a relation r, choose its (direct) parent $r' \in \mathcal{L}$.
  - A $|\mathcal{L}|$-way multi-class classification problem.

h	t
(Apple,	Jobs)
(Microsoft,	Gates)
(Facebook,	Zuckerberg)
…	

r=founded by

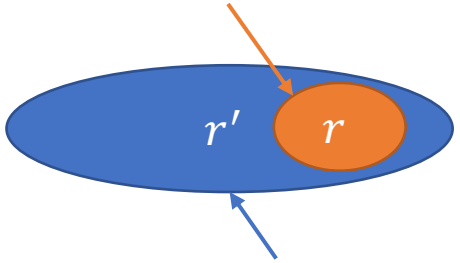

r'=creator

h	t
(Apple,	Jobs)
(Microsoft,	Gates)
(Facebook,	Zuckerberg)
…	
(Anna Karenina,	Leo Tolstoy)
(Pride and Prejudice,	Jane Austen)
…	

Relation Entailment Task Definition

- Notations
  - Head and tail entities $h, t \in \mathcal{E}$.
  - Relations $r \in \mathcal{R}$.
  - Instances of a relation $C_r = \{(h, r, t)^{(i)}\}_i$.
- Relation entailment
  - $r \vDash r'$ if and only if $C_r \subseteq C_{r'}$.
- Task of predicting relation entailment
  - Given a relation r, choose its (direct) parent $r' \in \mathcal{L}$.
  - A $|\mathcal{L}|$-way multi-class classification problem.

h	t
(Apple,	Jobs)
(Microsoft,	Gates)
(Facebook,	Zuckerberg)
…	

r=founded by

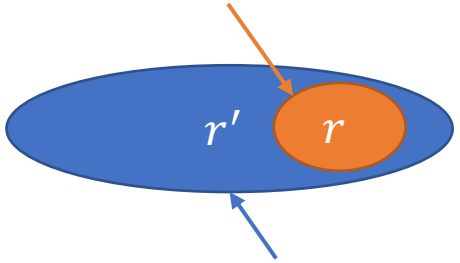

r'=creator

h	t
(Apple,	Jobs)
(Microsoft,	Gates)
(Facebook,	Zuckerberg)
…	
(Anna Karenina,	Leo Tolstoy)
(Pride and Prejudice,	Jane Austen)
…	

RelEnt Dataset

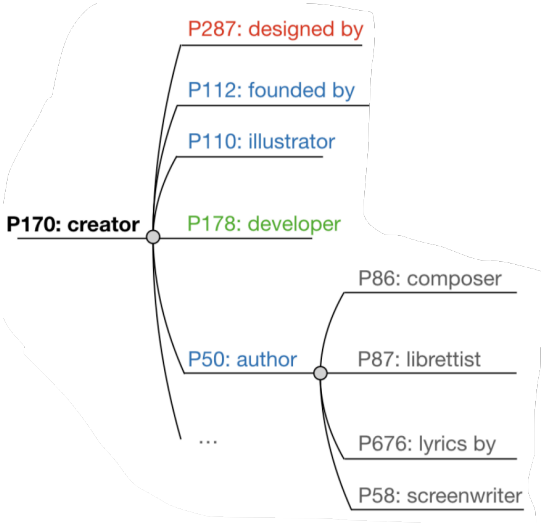

RelEnt Dataset

1. Instances collection

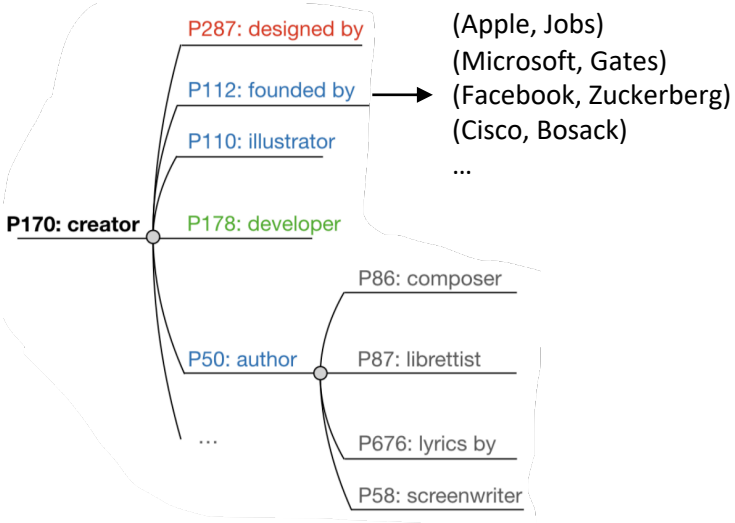

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
(Cisco, Bosack)
…

RelEnt Dataset

1. Instances collection

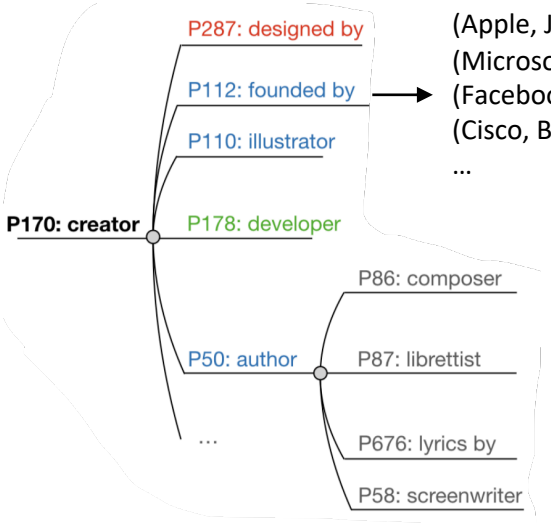

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
(Cisco, Bosack)
…

2. Downsampling

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
…

RelEnt Dataset

1. Instances collection

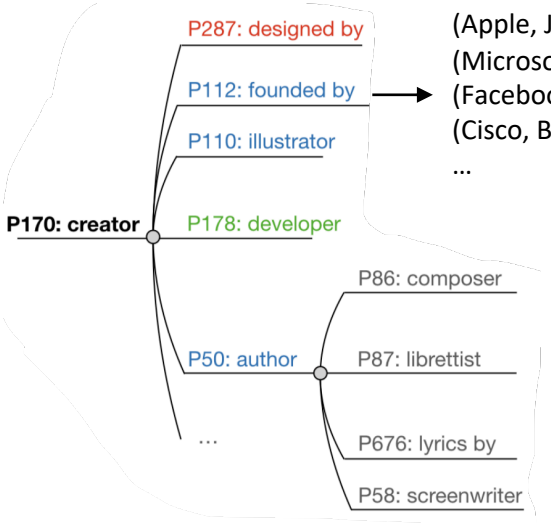

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
(Cisco, Bosack)
...

2. Downsampling

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
...

3. Relation expansion

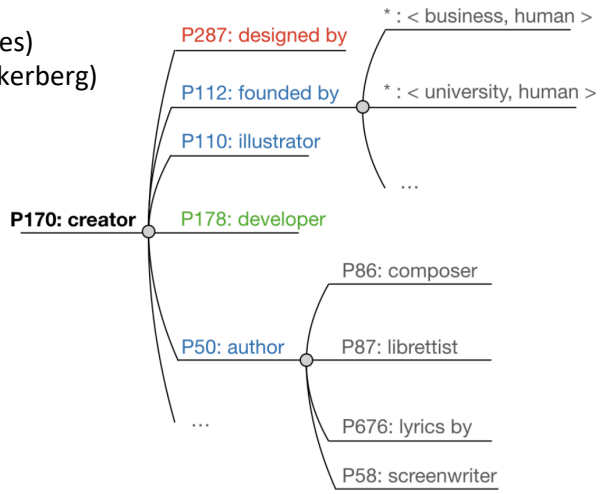

RelEnt Dataset

1. Instances collection **2. Downsampling** **3. Relation expansion**

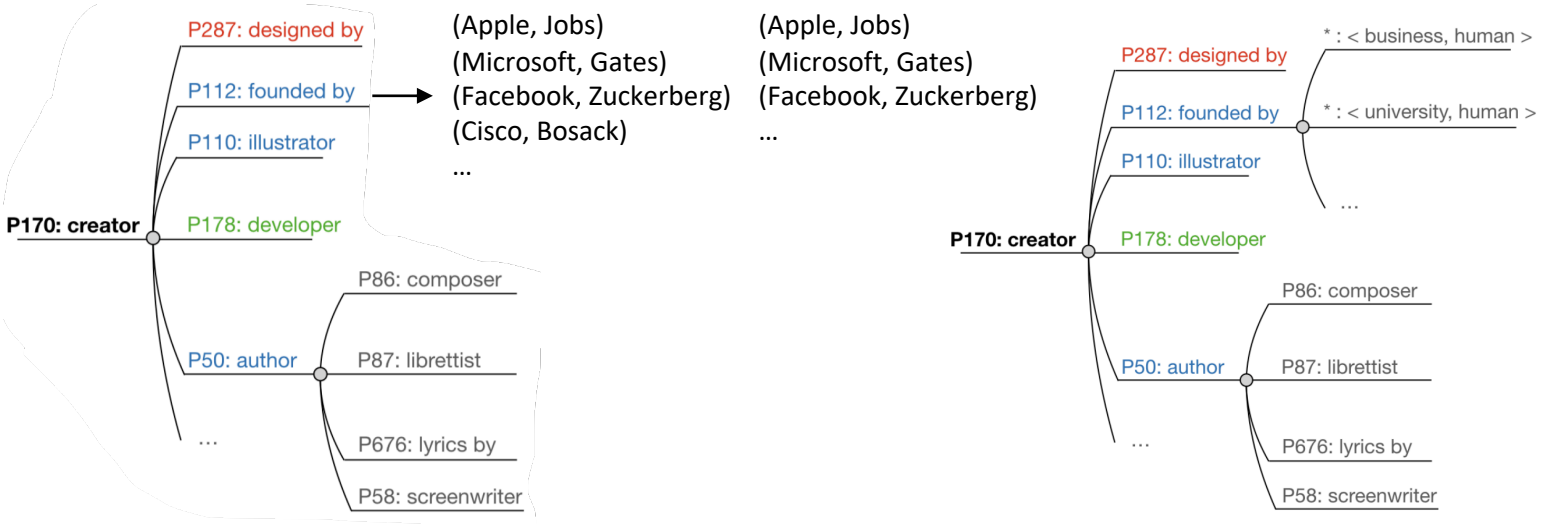

parent	Sub-relations
parent organization	<laboratory, university>, <airline, airline>, <record label, record label>, …
architectural style	<railway station, architectural style>, <church, architectural style>, ...
award received	<film, Academy Awards>, <human, campaign medal>, <human, scholarship>, ...

RelEnt Dataset

1. Instances collection **2. Downsampling** **3. Relation expansion** **4. Entity linking**

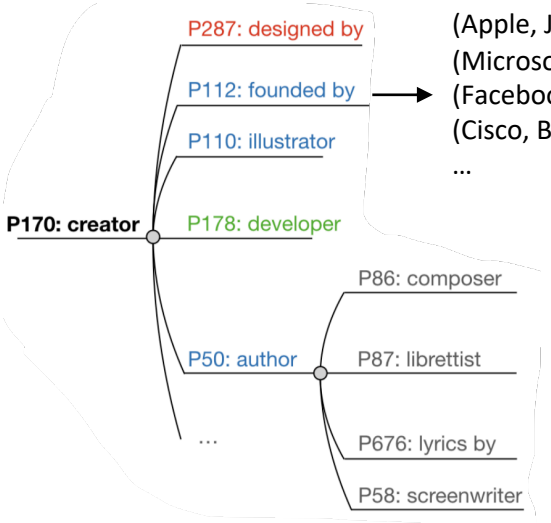

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
(Cisco, Bosack)
...

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
...

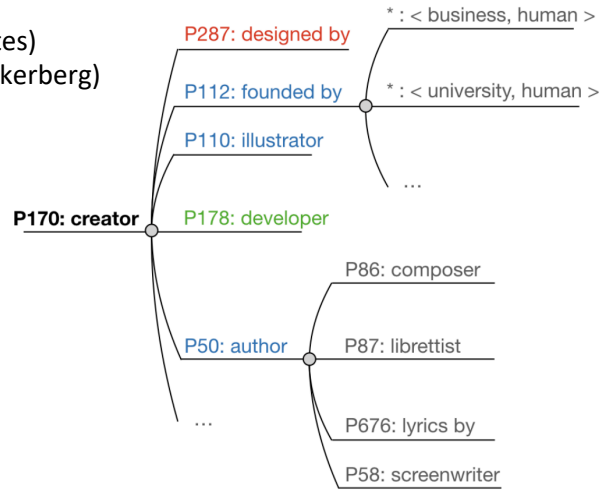

WIKIPEDIA
The Free Encyclopedia

parent	Sub-relations
parent organization	<laboratory, university>, <airline, airline>, <record label, record label>, …
architectural style	<railway station, architectural style>, <church, architectural style>, ...
award received	<film, Academy Awards>, <human, campaign medal>, <human, scholarship>, ...

RelEnt Dataset

1. Instances collection

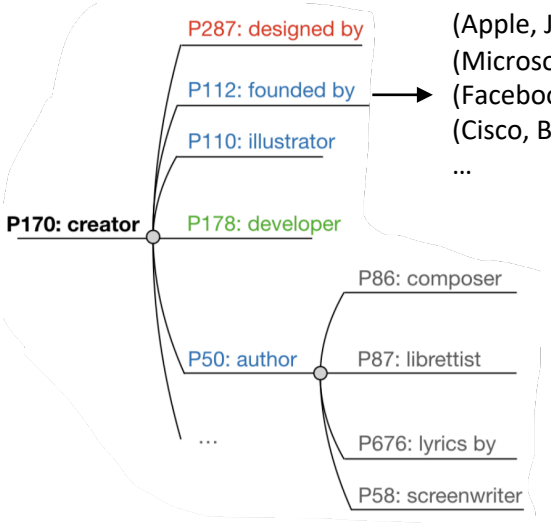

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
(Cisco, Bosack)
…

2. Downsampling

(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
…

3. Relation expansion

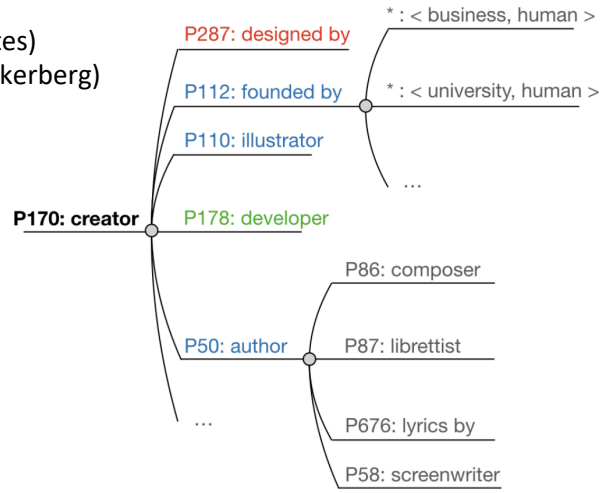

4. Entity linking

WIKIPEDIA
The Free Encyclopedia

5. Train/dev/test split

Train
founded by ⊨ creator
illustrator ⊨ creator
author ⊨ creator
Dev
developer ⊨ creator
Test
designed by ⊨ creator

parent	Sub-relations
parent organization	<laboratory, university>, <airline, airline>, <record label, record label>, …
architectural style	<railway station, architectural style>, <church, architectural style>, ...
award received	<film, Academy Awards>, <human, campaign medal>, <human, scholarship>, ...

Relation Representation

- With structured information

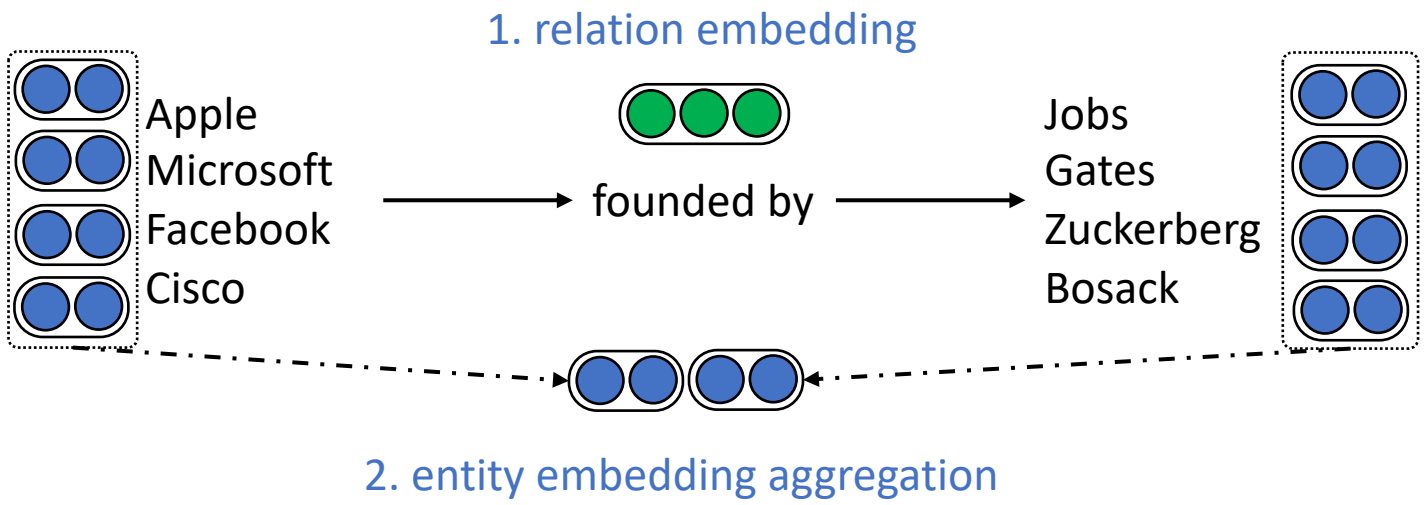

Relation Representation

- With textual information

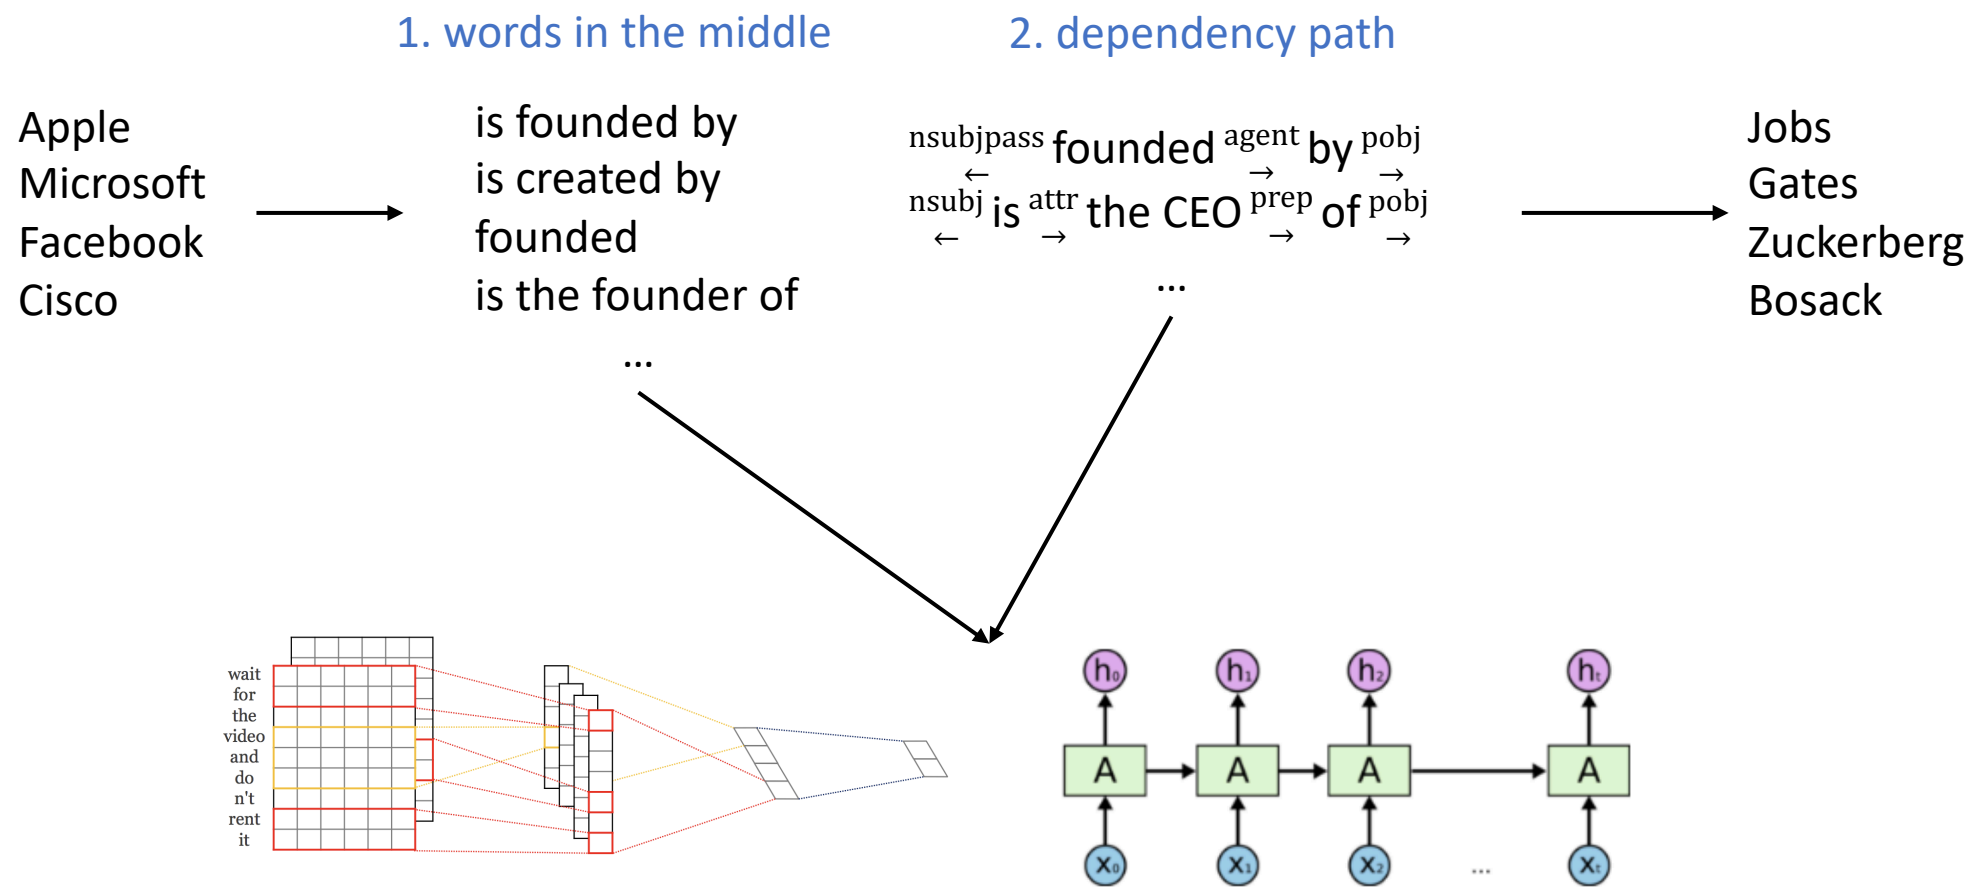

Relation Representation

- Distribution-based

author
(Pride and Prejudice, Jane Austen)
(The Lord of the Rings, J. R. R. Tolkien)
(Anna Karenina, Leo Tolstoy)
…

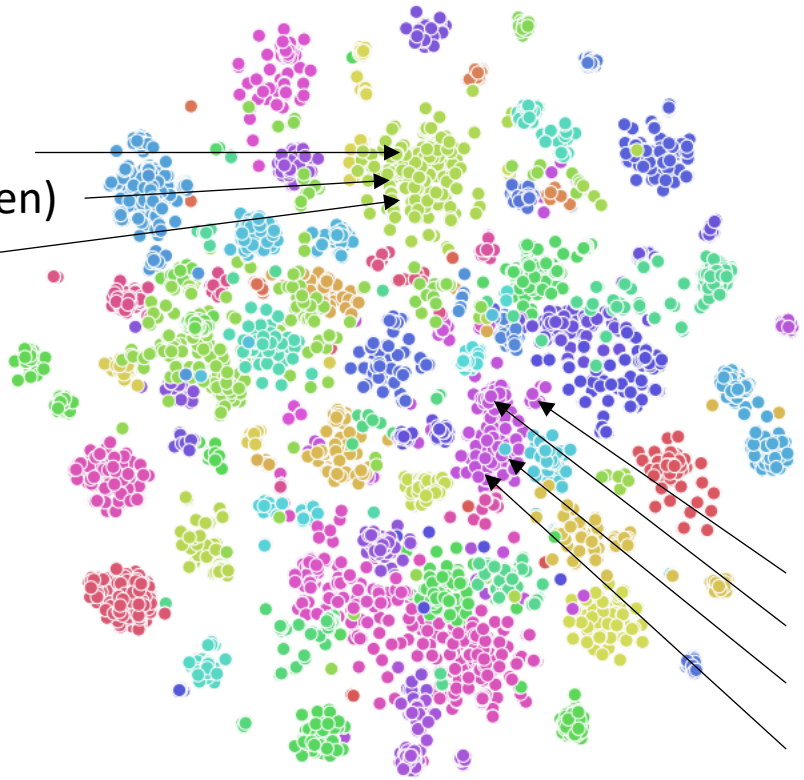

founded by
(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
(Cisco, Bosack)
…

Relation Representation

- Distribution-based

author
(Pride and Prejudice, Jane Austen)
(The Lord of the Rings, J. R. R. Tolkien)
(Anna Karenina, Leo Tolstoy)
…

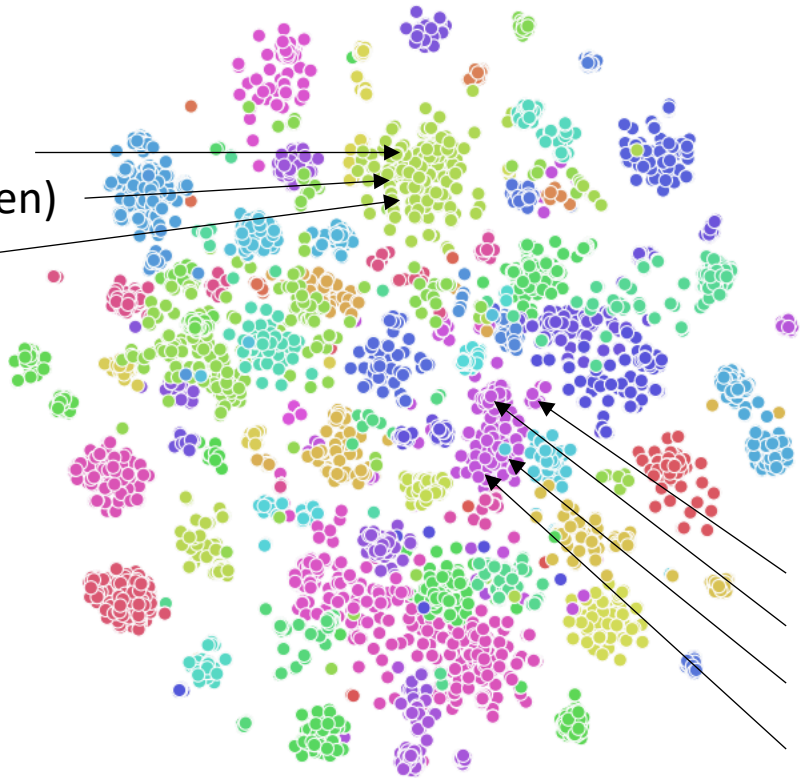

founded by
(Apple, Jobs)
(Microsoft, Gates)
(Facebook, Zuckerberg)
(Cisco, Bosack)
…

Kernel density estimation with a Gaussian kernel

Relation Entailment Prediction

Unsupervised methods

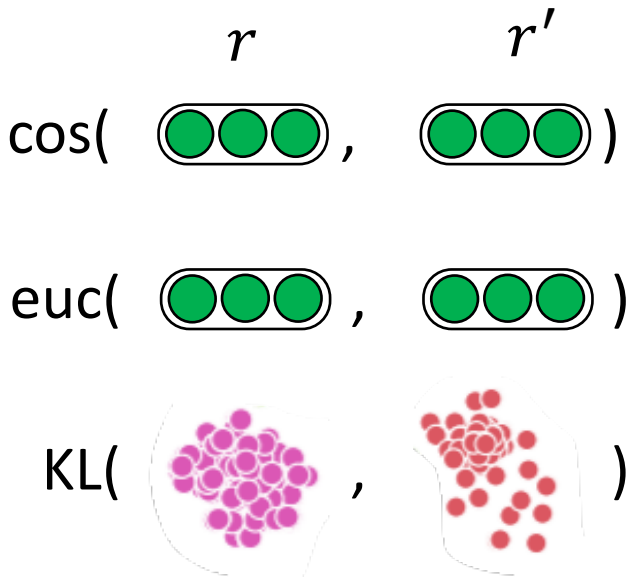

Supervised methods

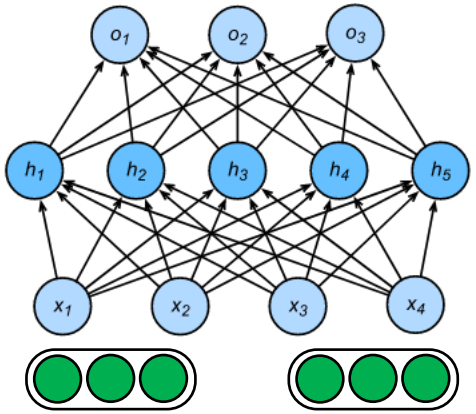

Experimental Settings

- RelEnt Dataset
  - #Train, #Dev., #Test relations: 2055, 804, 692
  - #Classes: 498

- Evaluation Metrics
  - Accuracy@1, Accuracy@3, and mean reciprocal rank (MRR)

Unsupervised Methods' Results

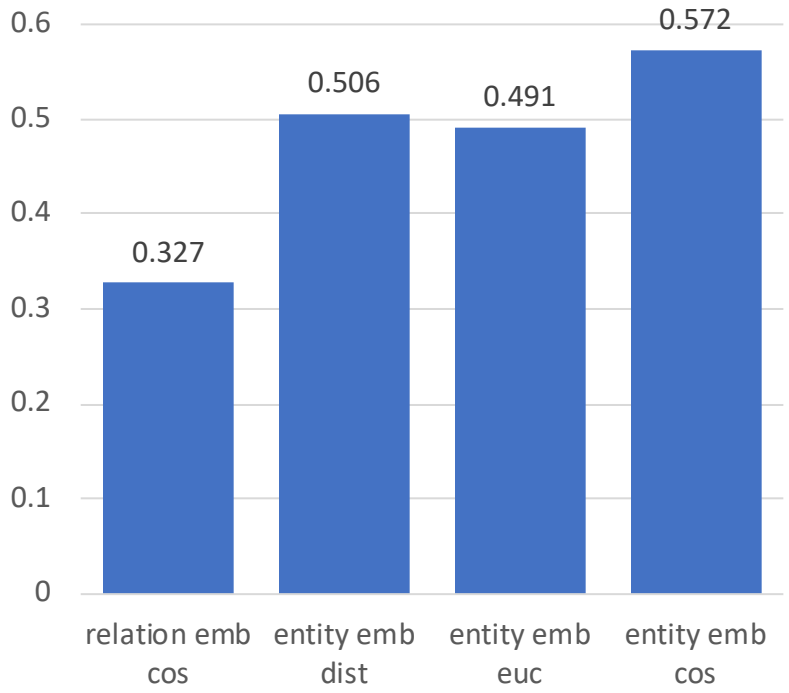

Acc@1 of different unsupervised methods with TransE.

Unsupervised Methods' Results

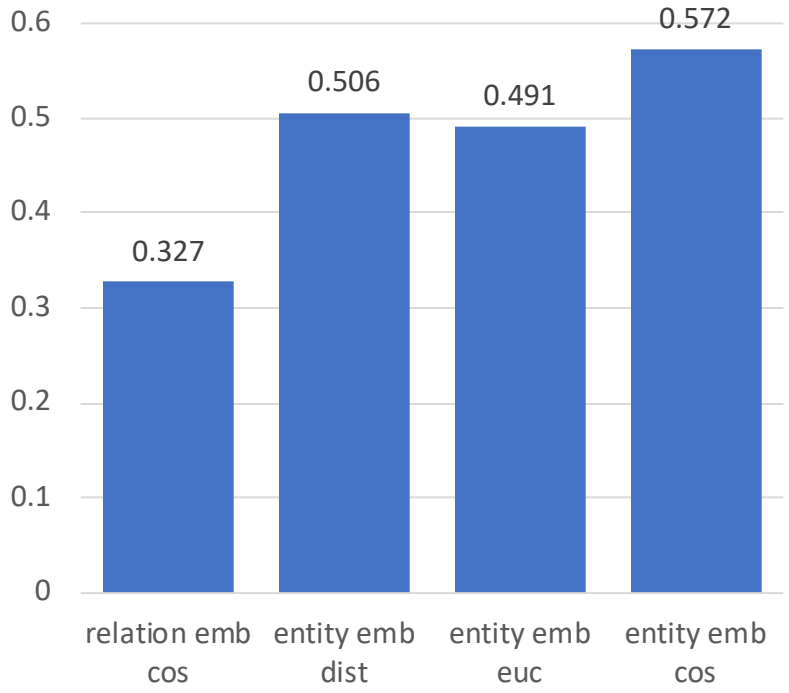

Acc@1 of different unsupervised methods with TransE.

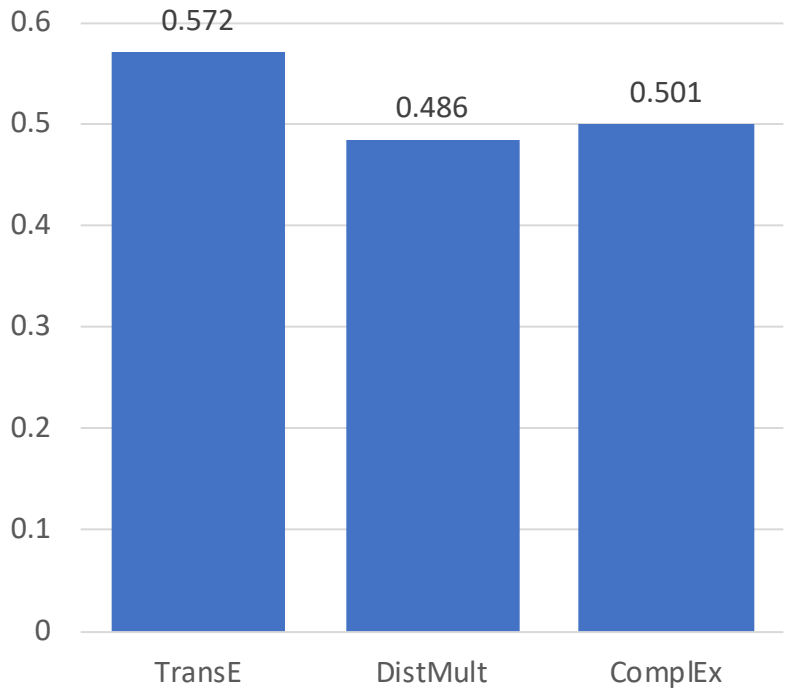

Acc@1 of entity embedding with cosine using different KG representations.

Supervised Methods' Results

- Supervised > unsupervised.
- Textual information is complementary to structured information.

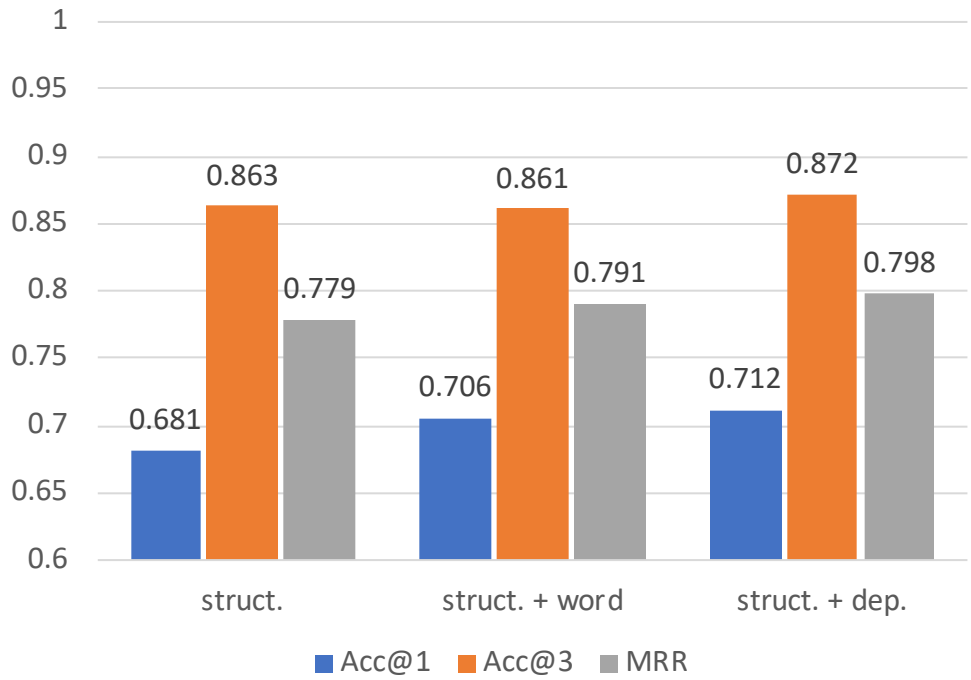

Error cases

Parent	Child (train)	Child (test)
follows	has cause	replaces
instance of	taxon rank	legal form
participant	performer	participating team

(2010 Wimbledon Championships, Roger Federer)
(First Continental Congress, George Washington)
(Hambach Festival, Ludwig Börne)

(Runaway, Linkin Park)
(The Freewheelin' Bob Dylan, Bob Dylan)

(1977 UEFA Cup Final, Juventus FC)
(2016–17 Premier League, Watford F.C.)
(1956 Wrestling World Cup, Iran)

Carnegie Mellon University School of Computer Science

BOSCH Invented for life

The Ohio State University Department of Computer Science and Engineering

NEULAB

Take away

1. Both structured and textual information contribute to relation entailment prediction.

2. Relation entailment prediction requires high-level abstraction.

Paper: https://openreview.net/pdf?id=ToTf_MX7Vn
Code: https://github.com/jzbjyb/RelEnt