



Demystifying Relational Latent Representations

Sebastijan Dumančić, Hendrik Blockeel

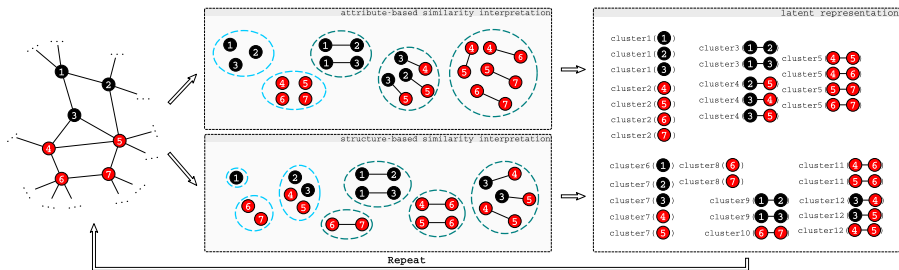
DTAI, KU Leuven

September 6, ILP 2017

- 1 Introduction
- 2 Understanding latent features
- 3 Properties of latent spaces

1 – Learning Relational Latent Representations 3/24

Learning versatile relational latent features with clustering and variety of similarities (CUR²LED)



[Dumančić and Blockeel, IJCAI 2017]

Benefits:

- better performance
- simpler models


[with some overhead]

Questions to be answered

- ① *Can we interpret latent features?*
 - (approximate) definition of latent features/relations
- ② *What makes them effective?*
 - distinctive properties?


- 1 Introduction
- 2 Understanding latent features
- 3 Properties of latent spaces

latent features = clusters of vertices (instances) and edges (relationships)




key idea: cluster prototype represents the meaning of a feature

CUR²LED uses ReCeNT as a similarity measure for relational data
⇒ views instances as **neighbourhood trees**



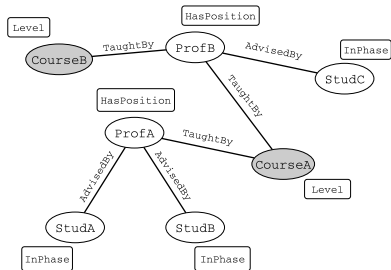
Data




Neighbourhood tree

[Dumančić and Blockeel, MLJ 2017]

CUR²LED uses ReCeNT as a similarity measure for relational data
 ⇒ views instances as **neighbourhood trees**



Data




Neighbourhood tree

key idea: *mean tree* represents the meaning of a feature

CUR²LED requires a (set of) **similarity interpretation(s)**

→ a specification what similarity reflects


- attribute sim
- neighbourhood attributes sim
- neighbourhood identity
- edge labels
- connectedness



key idea: find a mean tree given the similarity interpretation

CUR²LED compares neighbourhood trees by comparing distributions of elements within them


- elements selected by the similarity interpretation
- attributes values, edge types, identities ...



key idea: mean tree \approx elements that appear in all NTs (in a cluster) with similar frequency


Given a set of neighbourhood tree and a similarity interpretation ...

1.






Calculate the relative frequencies of elements within a tree

1.




2.



Summarize the relative frequencies of unique elements across trees

1.



2.

advisedBy: $\frac{2}{3}$
teaches: $\frac{1}{3}$

advised: $\frac{1}{3}$
member: $\frac{1}{3}$
teaches: $\frac{1}{3}$


advisedBy: $\frac{2}{3}$
teaches: $\frac{1}{3}$

3.




```
advisedBy -> mean: 0.55  
              std: 0.15  
teaches -> mean: 0.33  
              std: 0.0  
member -> mean: 0.11  
              std: 0.15
```

Select elements with low standard deviation

1.



2.




3.

```


advisedBy -> mean: 0.55
              std: 0.15
teaches -> mean: 0.33
              std: 0.0
member -> mean: 0.11
              std: 0.15
  
```

(θ -confidence) An element with mean value μ and standard deviation σ in a cluster, is said to be θ -confident if $\sigma < \theta \cdot \mu$.

Use case: IMDB

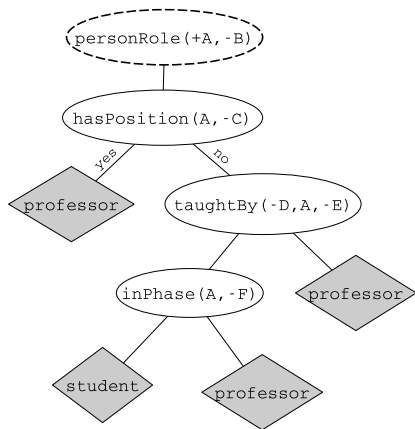


a)




b)

Use case: UWCSE



a)



b)

- 1 Introduction
- 2 Understanding latent features
- 3 Properties of latent spaces

Properties of latent spaces:

- **label entropy**

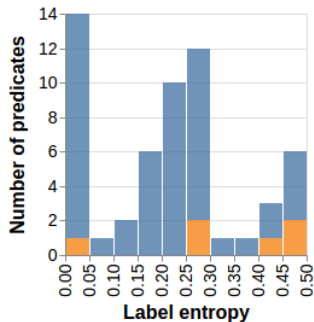
- distribution of labels within true instantiations of predicates
- proxy to a quantification of learning difficulty

- **sparsity**


- modelling *local* vs. *global*
- concept spread across a small number of local regions is easier to capture

- **redundancy**

- CUR²LED creates many features - are all of them necessary?

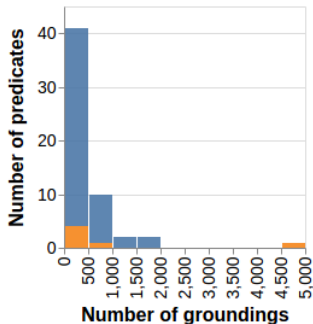


improved performance (3)

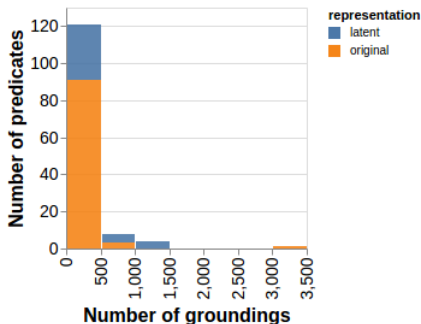


no improvement (1)

when performance increases, latent representation has many predicates of low label entropy



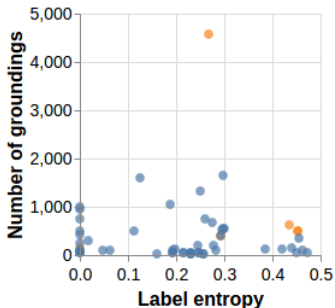
improved performance (3)




no improvement (1)

when performance increases, latent representation is sparser than the original one

Trivial explanation: many predicates with a very small number of true instantiations (not helpful)



improved performance (3)



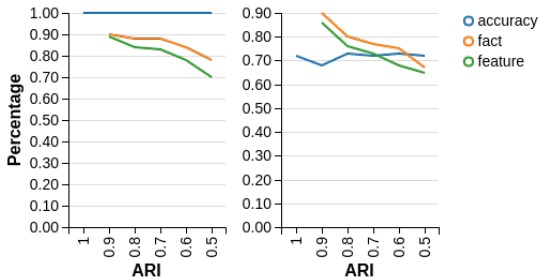
no improvement (1)

... not what's happening here: latent predicates have a comparable number of true instantiations

latent features successfully identify local regions in the instance space that match well with the provided labels

CUR²LED creates a lot of features

- similarity interpretations considered independently
- but many instances might be identical in several similarity interpretations
- every time a new clustering is obtained, check how much it overlaps with an existing clusterings using the adjusted Rand index



Thank you!
Questions?