

KR 2023

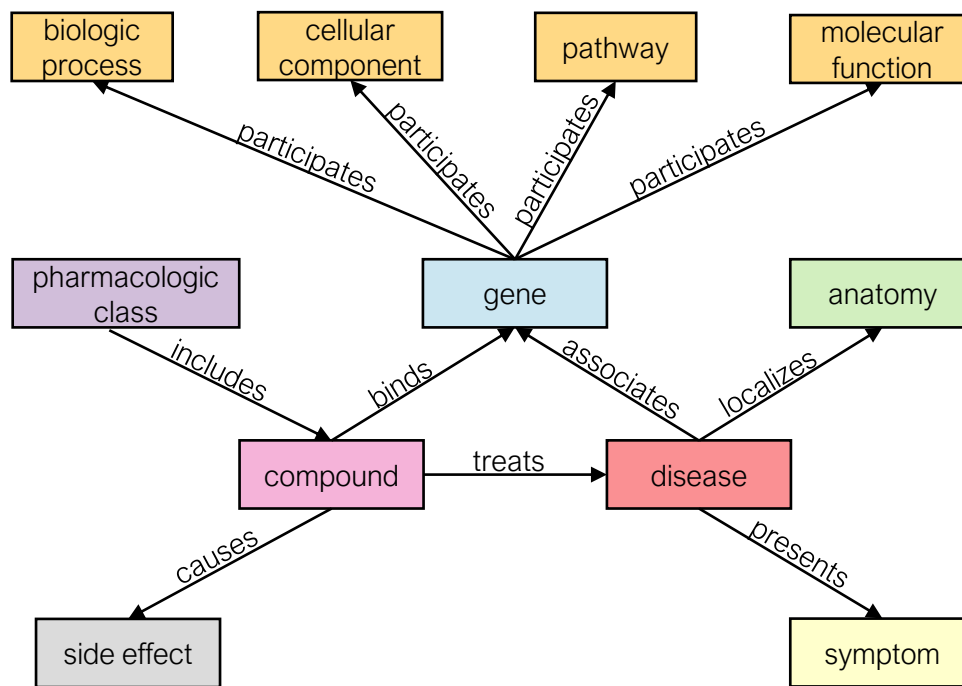
# EXPLAINABLE REPRESENTATIONS FOR RELATION PREDICTION IN KNOWLEDGE GRAPHS

Rita T. Sousa, Sara Silva, Catia Pesquita  
LASIGE, Faculdade de Ciências da Universidade de Lisboa

September 7, 2023

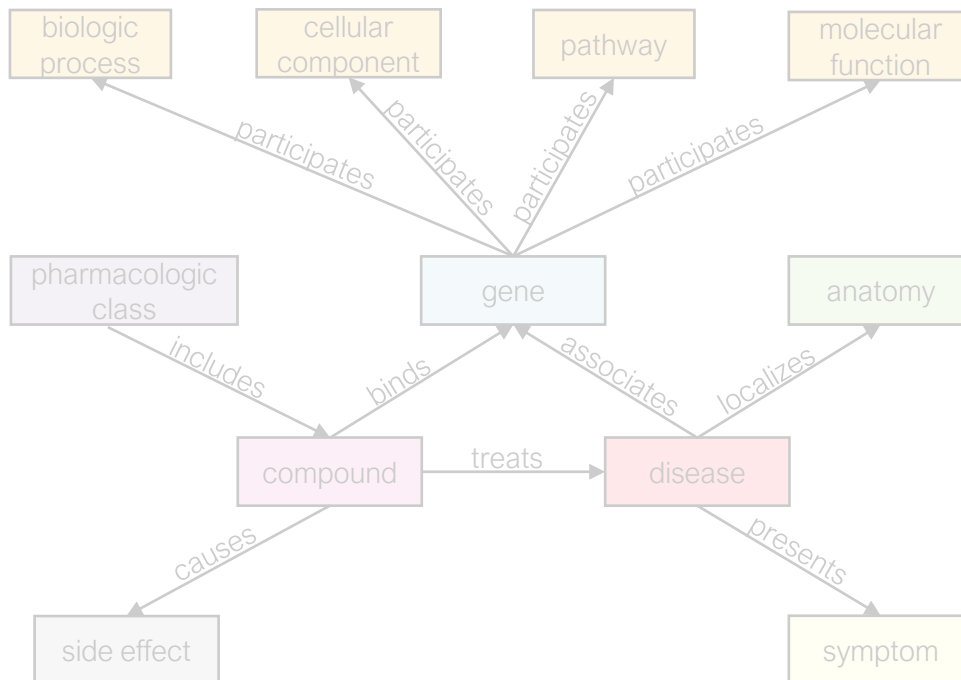
20th International Conference on Principles of Knowledge Representation and Reasoning

**Knowledge graphs (KGs) represent real-world entities and their relations in a semantically-rich structure, allowing the description of complex natural phenomena.**



KGs are explainable by design.

## However, most machine learning methods that work over KGs generate vector representations that are non-explainable.

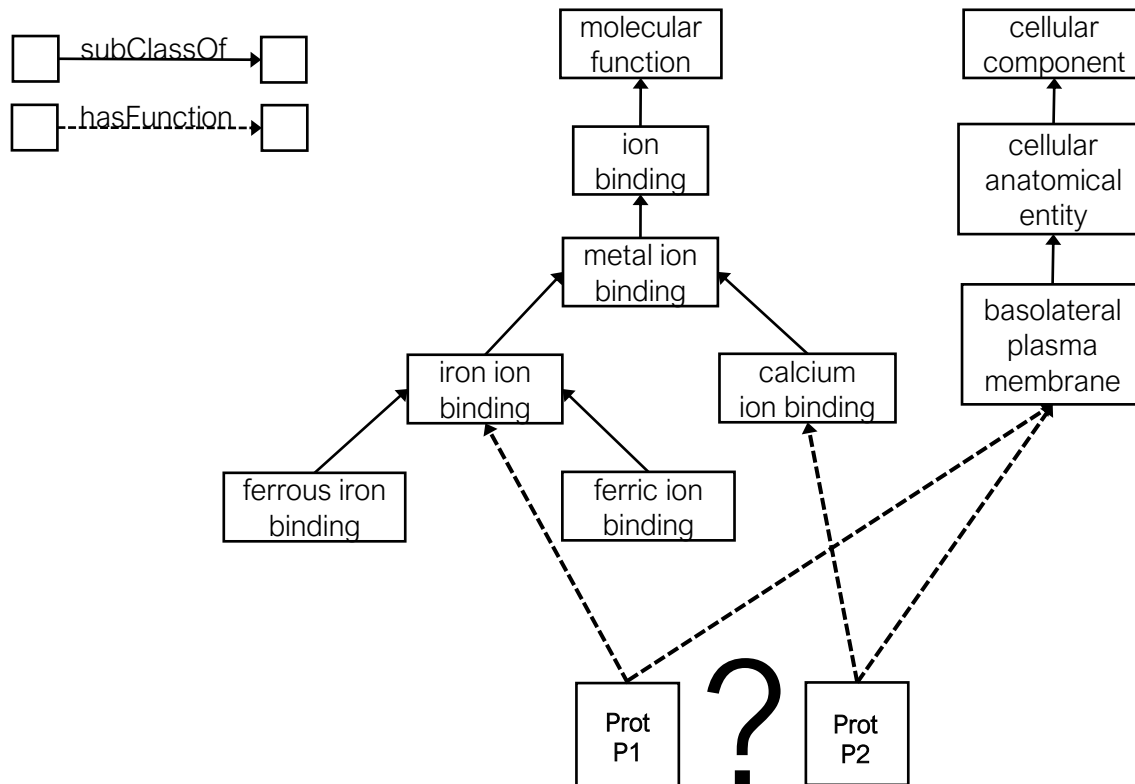


biologic process	0.6	0.8	...
cellular component	0.3	0.2	...
pathway	0.1	0.7	...
molecular function	0.7	0.6	...
pharmacologic class	0.5	0.5	...
gene	0.4	0.2	...

compound	0.1	0.9	...
disease	0.4	0.4	...
symptom	0.2	0.8	...
side effect	0.7	0.1	...
anatomy	0.2	0.6	...

**Can we use these KGs without  
sacrificing the full and rich  
explainability offered by them?**

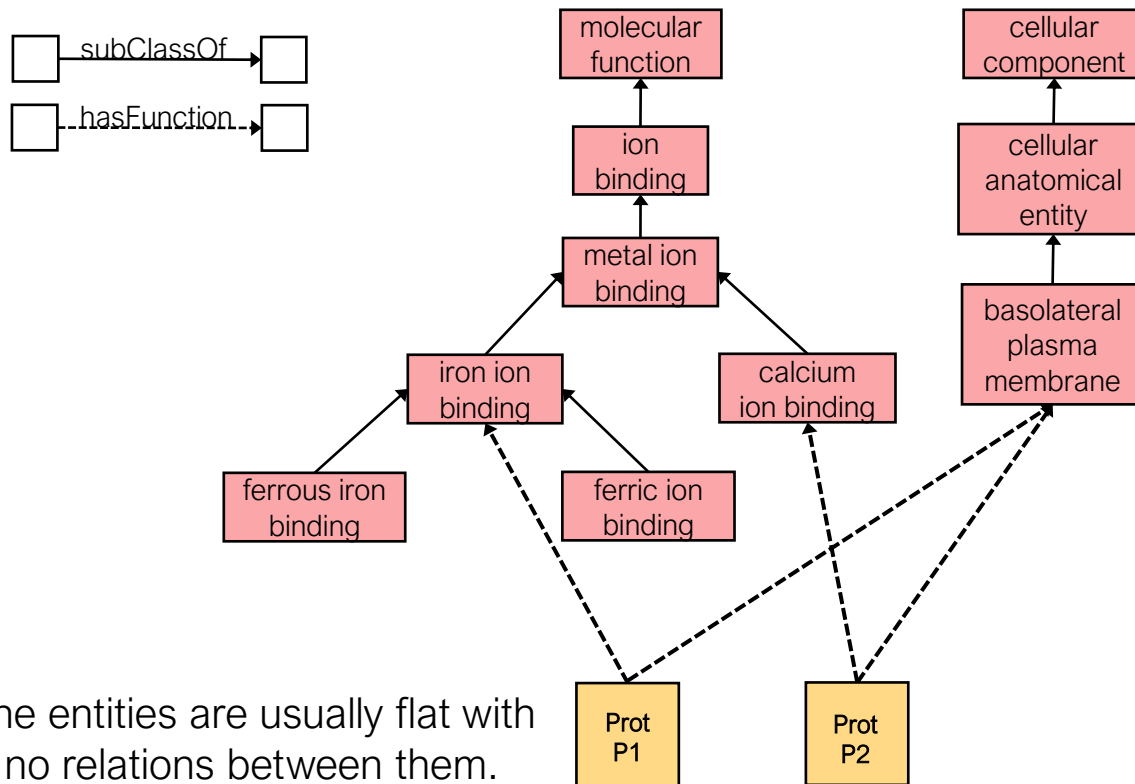
## In relation prediction, understanding which features better explain a relation is crucial to support critical applications.



Relation prediction is an important task in the biomedical domain:

- Protein-protein interaction (PPI) prediction;
- Gene-disease association (GDA) prediction;
- Drug-Drug interaction (DDI) prediction;
- Drug-target association (DTI) prediction.

## Ontology-rich KGs use ontologies to describe entities.



The entities are usually flat with no relations between them.

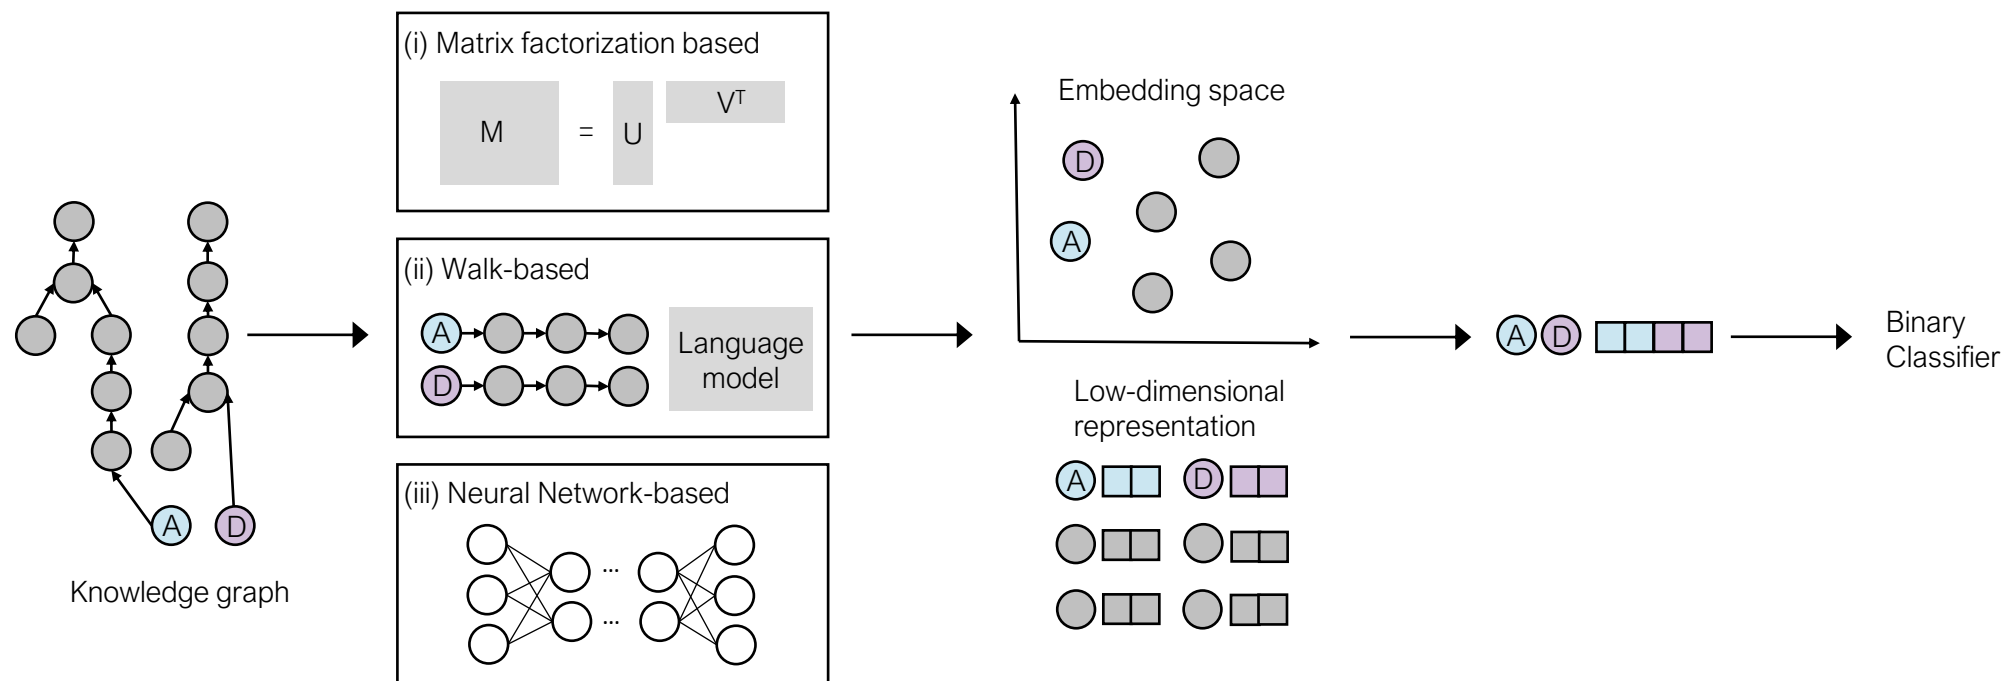
Two types of vertices:

- individual entities;
- ontology classes.

Two types of edges:

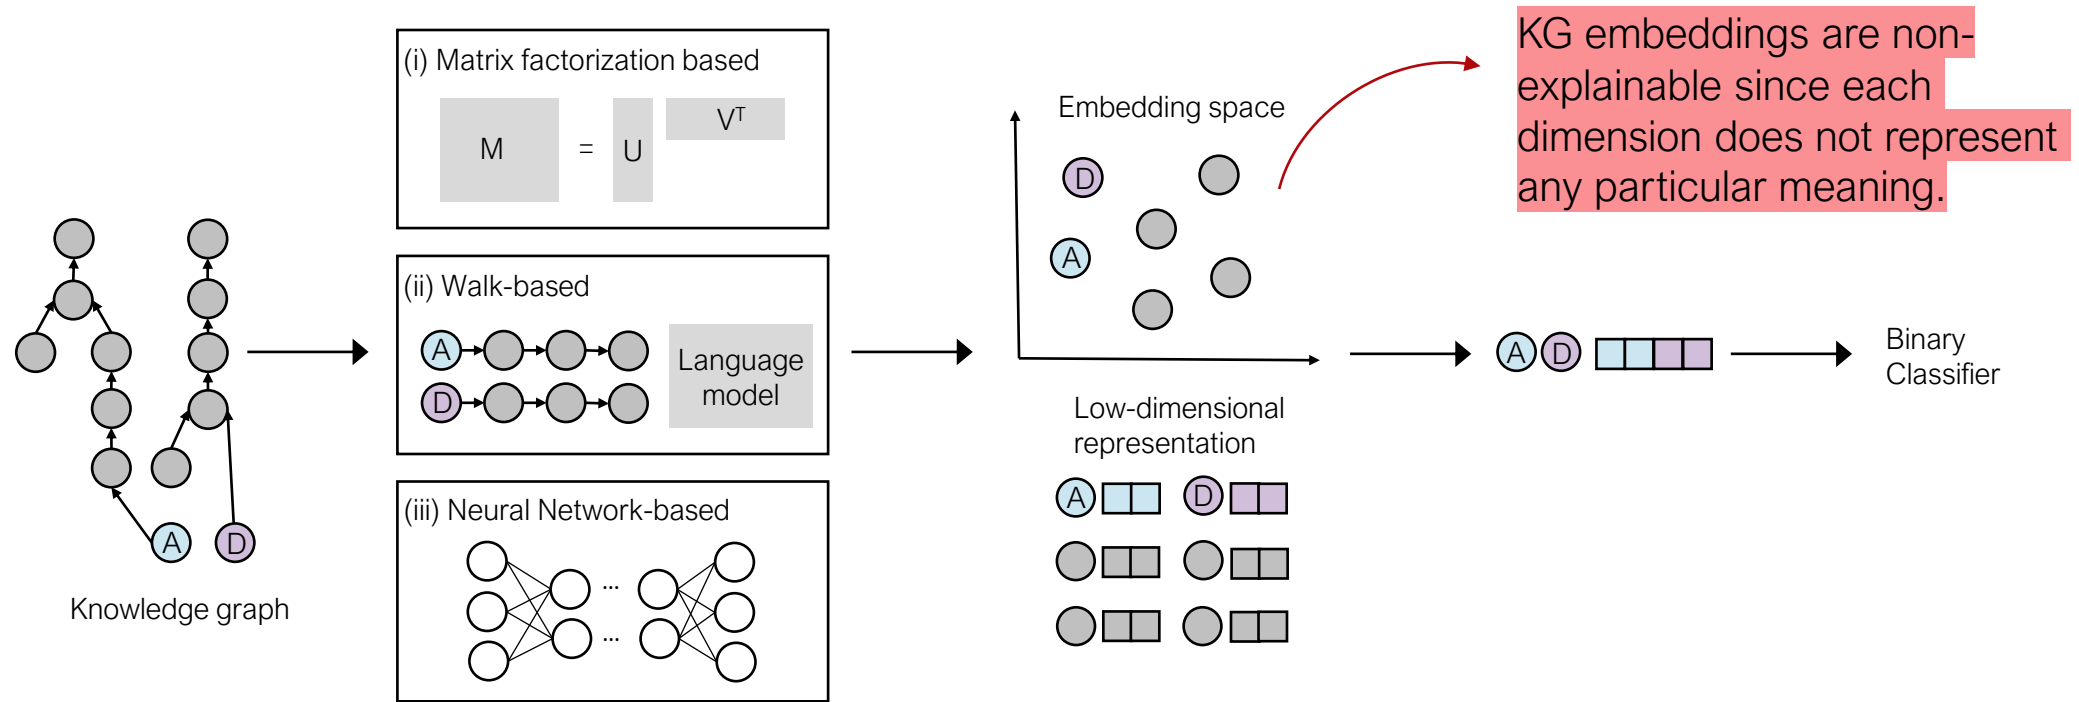
- relate ontology classes to each other ( $\longrightarrow$ );
- link entities to the classes that describe them ( $\dashrightarrow$ ).

**Most relation prediction approaches rely on KG embeddings to bridge the gap between the complex representations a KG affords and the vectorial representations most machine learning methods take as input.**



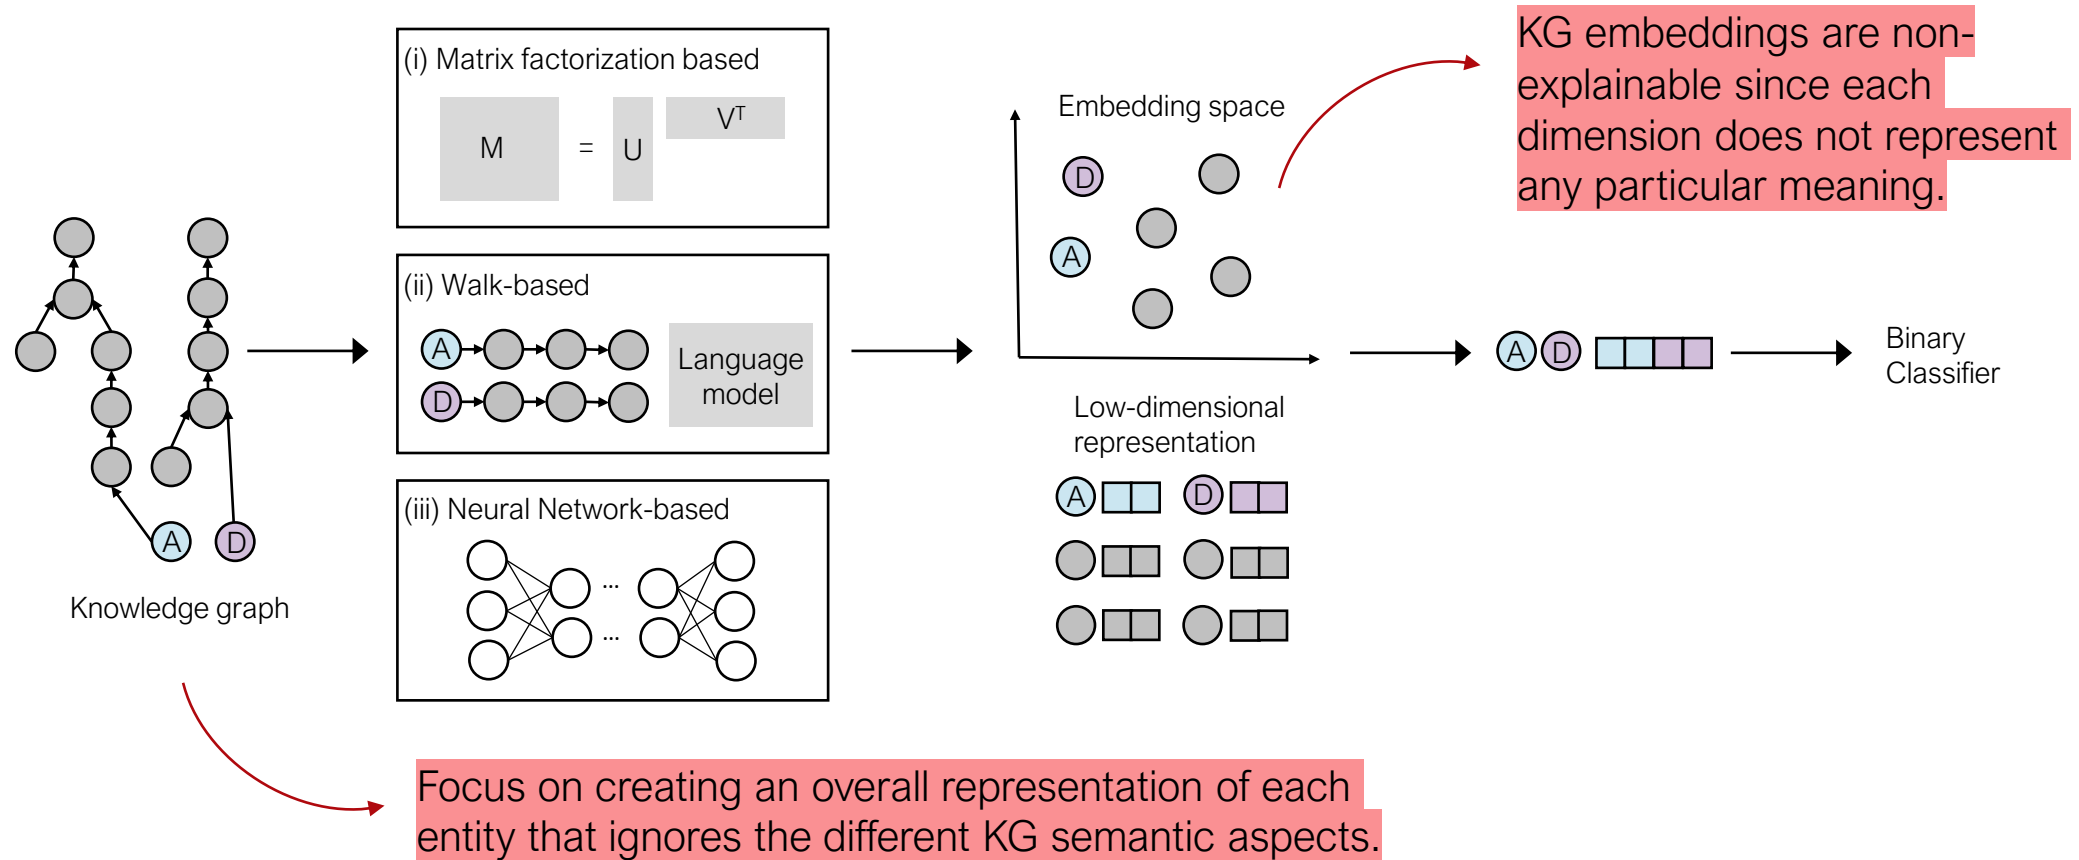
There are several types of KG embeddings, including translational models, semantic matching models, or random walk-based KG embedding approaches.

# The KG embedding-based approaches generate non-explainable predictions, which hinders their application in relation prediction.

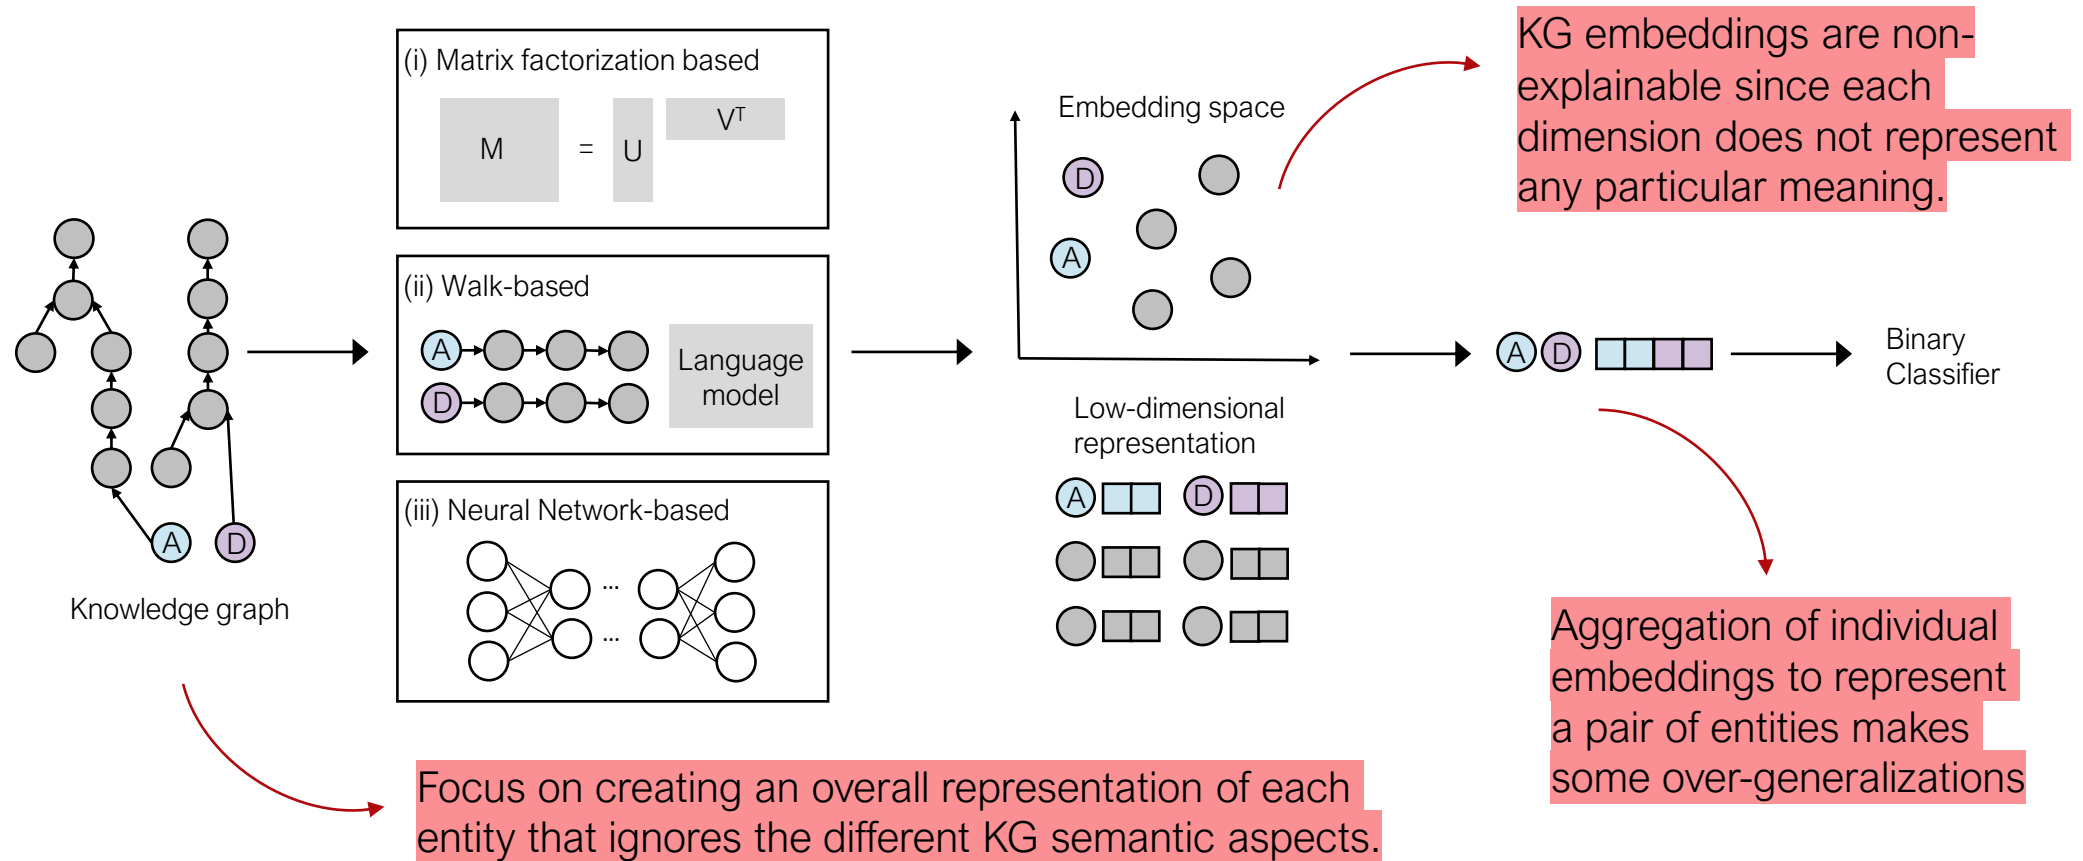





# The KG embedding-based approaches generate non-explainable predictions, which hinders their application in relation prediction.



# The KG embedding-based approaches generate non-explainable predictions, which hinders their application in relation prediction.




# PROBLEM OVERVIEW



- i. Generation of latent representations that represent an entity pair directly.
- ii. Generation of latent representations that are amenable to explanation and can capture the relevant semantic aspects for relation prediction.

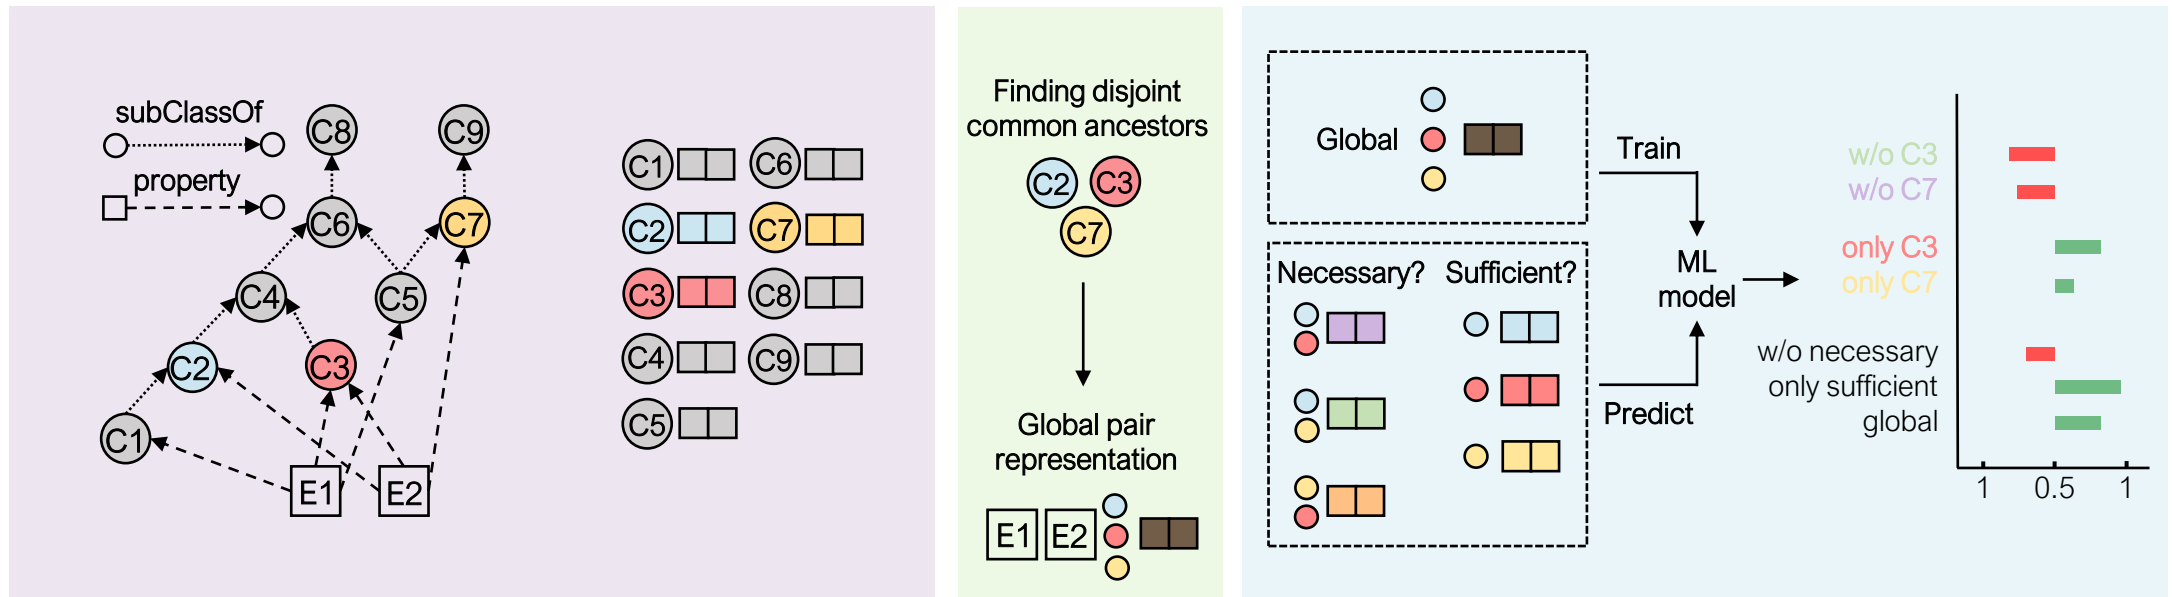
Low-dimensional representation





# SEEK is a novel approach for explainable representations to support relation prediction in KGs.

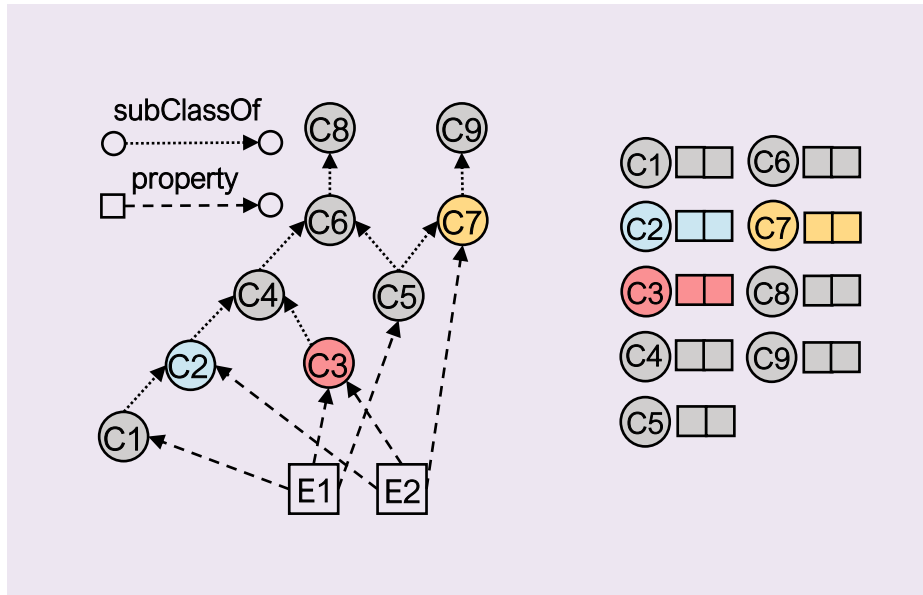
SEEK generates perturbed representations to identify the relevant semantic aspects of the KG that explain a relation.



# SEEK



## (I) Generating the RDF Graph and Learning Embeddings



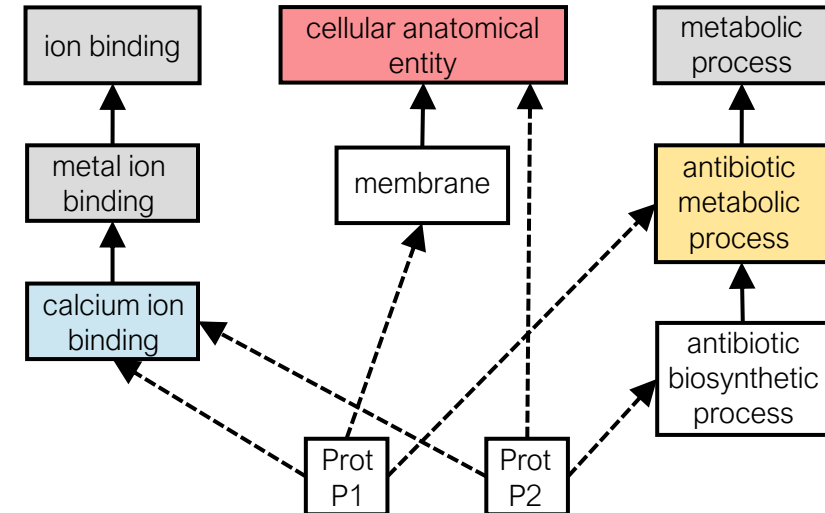

- The majority of graph processing and analysis tools require RDF graphs.
- The conversion of the KG involves transforming axioms into RDF triples.
- After conversion, a KG embedding method is employed to learn latent representation of all the ontology classes in the KG.

# SEEK



## (II) Finding Shared Semantic Aspects and Generating Pair Representations

- A pair is represented by the set of semantic aspects they share.
- The shared semantic aspects is the set of disjoint common ancestors computed over the set of classes that describe each entity.
- The pair representation is computed by aggregating the embeddings of each shared semantic aspect.



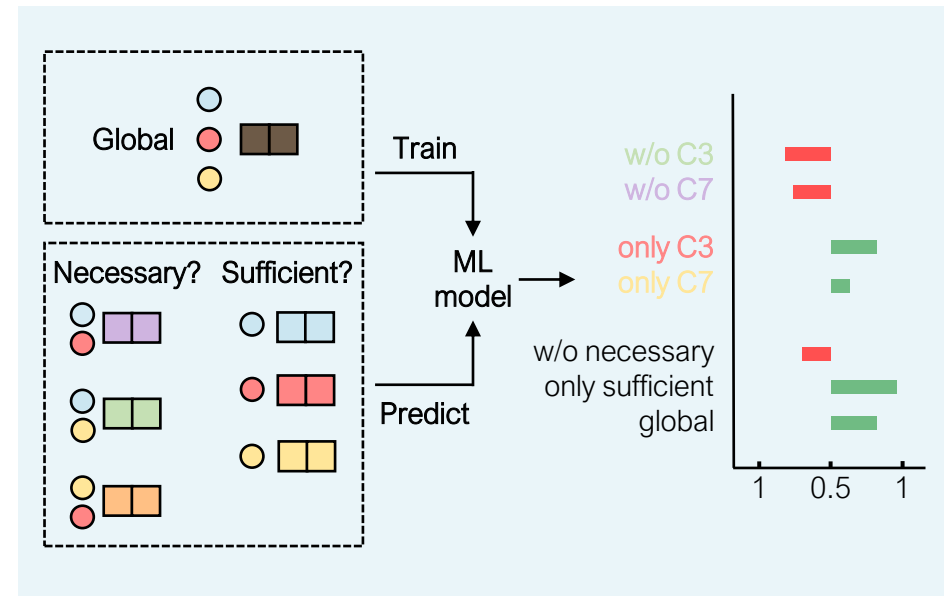
The shared semantic aspects of proteins P1 and P2 are: 'calcium ion binding', 'cellular anatomical entity', and 'antibiotic metabolic process'.

# SEEK




## (III) Predicting and Explaining

- The model is trained using the global representation. To generate explanations, multiple representations are generated that differ by the presence or absence of a semantic aspect. **Does not require retraining the embeddings.**
- A **necessary explanation** is a shared semantic aspect that, when removed from the pair representation, causes the classifier to change its prediction.
- A **sufficient explanation** is a shared semantic aspect that, when used alone to represent a pair, causes the classifier to maintain its prediction.



# SEEK EXPLANATION

An explanation is the set of the most relevant shared semantic aspects identified as necessary or sufficient. The final explanation can be represented as a chart.






# EXPERIMENTAL SETUP

SEEK is evaluated on two biomedical relation prediction tasks:

## PPI prediction

- Target relations from STRING.
- Gene Ontology KG used to describe proteins.



- Known interactions:
- from curated databases
  - experimentally determined
- Predicted interactions:
- gene neighborhood
  - gene fusions
  - gene co-occurrence
- Others:
- text mining
  - co-expression
  - protein homology

Figure adapted from STRING database.

## GDA prediction

- Target relations from DisGeNET.
- Human Phenotype Ontology KG describe genes and diseases.

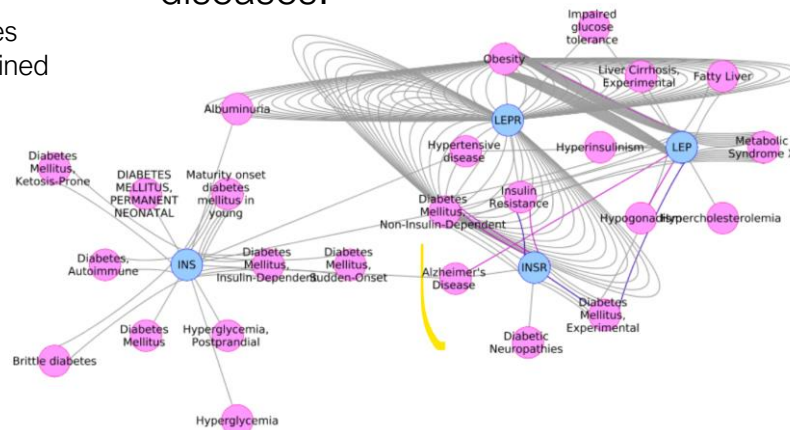


Figure extracted from DisGeNET-app.

	PPI	GDA
Ontology classes	50 422	15 656
Literals and blank nodes	462 874	443 489
Instances	6 738	4 523
Annotations	349 500	160 009
Positive Pairs	23 571	8 189
Negative Pairs	23 571	8 189

## PERFORMANCE EVALUATION

Median of f1-score comparing SEEK to the baseline when coupled with different supervised machine learning methods.\*

### KG embedding methods

- RDF2Vec;
- OWL2Vec\*;
- TransE;
- TransH;
- distMult.

### Supervised ML models

- Random Forest (RF);
- eXtreme Gradient Boosting (XGB);
- Multilayer Perceptron (MLP).

Method	PPI Prediction		GDA Prediction		
	Baseline	SEEK	Baseline	SEEK	
RDF2Vec	XGB	0.910	<u>0.915</u>	<u>0.724</u>	0.719
	RF	0.902	<u>0.910</u>	<u>0.740</u>	0.723
	MLP	0.902	<u>0.917</u>	0.696	<u>0.703</u>
OWL2Vec*	XGB	0.888	<u>0.929</u>	0.688	<u>0.728</u>
	RF	0.875	<u>0.919</u>	0.690	<u>0.737</u>
	MLP	0.869	<u>0.931</u>	0.650	<u>0.720</u>
distMult	XGB	0.902	<u>0.912</u>	0.704	<u>0.722</u>
	RF	0.884	<u>0.905</u>	0.706	<u>0.716</u>
	MLP	<b>0.896</b>	0.888	<b>0.715</b>	0.698
TransE	XGB	0.637	<u>0.919</u>	0.510	<u>0.726</u>
	RF	0.583	<u>0.905</u>	0.502	<u>0.719</u>
	MLP	0.333	<u>0.920</u>	0.333	<u>0.711</u>
TransH	XGB	0.637	<u>0.919</u>	0.510	<u>0.726</u>
	RF	0.579	<u>0.910</u>	0.494	<u>0.720</u>
	MLP	0.333	<u>0.920</u>	0.333	<u>0.711</u>

\*f1-score values are underlined when improvements are statistically significant for the Wilcoxon test with p-value<0.05.

SEEK outperforms the baseline in all cases but one for PPI prediction, while achieving similar or improved scores for GDA.

## PERFORMANCE EVALUATION



t-SNE plots comparing SEEK to the baseline using RDF2Vec with **positive pairs in green** and **negative pairs in red**.

PPI prediction

GDA prediction



(a) Baseline

(b) SEEK



(a) Baseline

(b) SEEK




SEEK representations decrease the overlap between positive and negative pairs and, therefore, are likely to be capturing more meaningful representations.

## EXAMPLES OF EXPLANATIONS FOR PPI PREDICTION

True Positive

Paxilin and Integrin  $\alpha$ -4




Proteins involved in several diseases such as cancer.

Several aspects, such as focal adhesion, substrate adhesion-dependent cell spreading, cell migration and integrin binding, explain the interaction.

# EXAMPLES OF EXPLANATIONS FOR PPI PREDICTION

False Positive

Protransforming growth factor  $\alpha$  – Disks large homolog 2



Proteins involved in neurological injuries.

However, the literature describes interactions between proteins of the same family of the pair, indicating that this is likely a true but still unknown interaction.

Dlg5 interacts with the TGF- $\beta$  receptor and promotes its degradation

Edited by Gianni Cesareni

Takuhiro Sezaki <sup>a, b</sup>, Lucia Tomiyama <sup>a, b</sup>, Yasuhisa Kimura <sup>a</sup>, Kazumitsu Ueda <sup>a, b</sup>, Noriyuki Kioka <sup>a, b, c</sup>

## CONCLUSIONS

- **SEEK** is a novel approach for learning and explaining representations of KG entity pairs based on their shared semantic space for relation prediction.
- **SEEK** is evaluated on PPI prediction and GDA, two complex and core tasks in the biomedical domain.
- **SEEK** outperforms state-of-the-art learning representation methods in performance, while generating explanations that can identify critical factors driving biological phenomena.

## FUTURE WORK

- Investigate the minimal set of shared semantic aspects required to adequately explain a relation.
- Collaboration to test experimentally new interactions.



# Thank you for your attention.



<https://github.com/liseda-lab/seek>



risousa@ciencias.ulisboa.pt



@RitaTorresSousa

This work was funded by the FCT through LASIGE Research Unit (ref. UIDB/00408/2020 and ref. UIDP/00408/2020), and the FCT PhD grant (ref. SFRH/BD/145377/2019). It was also partially supported by the KATY project (No 101017453), and by HfPT: Health from Portugal under the Portuguese Plano de Recuperação e Resiliência.

# ADDITIONAL SLIDES



## EFFECTIVENESS OF EXPLANATIONS

Scenario where pairs are represented without the necessary shared semantic aspects. It is applied to correctly predicted relations. **The more negative  $\Delta$  is, the more effective.**

Scenario where pairs are represented by sufficient shared semantic aspects only. It is applied to incorrectly predicted relations. **A higher  $\Delta$  value indicates increased effectiveness.**

		PPI Prediction				GDA Prediction			
		RDF2Vec		OWL2Vec*		RDF2Vec		OWL2Vec*	
		MLP	XGB	MLP	XGB	MLP	XGB	MLP	XGB
Without necessary	$\Delta Pr$	-0.157	-0.109	-0.095	-0.099	-0.291	-0.296	-0.265	-0.332
	$\Delta Re$	-0.137	-0.120	-0.145	-0.131	-0.329	-0.220	-0.353	-0.208
	$\Delta F1$	-0.148	-0.113	-0.117	-0.113	-0.264	-0.225	-0.270	-0.256
Only sufficient	$\Delta Pr$	0.932	1.000	0.981	1.000	0.957	0.969	0.921	0.986
	$\Delta Re$	0.959	1.000	0.927	1.000	0.737	0.905	0.777	0.993
	$\Delta F1$	0.950	1.000	0.954	1.000	0.898	0.964	0.885	0.993

## EXPLANATION LENGTH

		PPI Prediction				GDA Prediction			
		RDF2Vec		OWL2Vec*		RDF2Vec		OWL2Vec*	
		Avg	Std	Avg	Std	Avg	Std	Avg	Std
sufficient	MLP	5.6	3.9	5.3	3.5	5.6	7.1	5.5	7.8
	XGB	6.2	3.9	6.3	4.1	6.0	8.3	6.0	9.4
	RF	5.6	3.7	5.9	3.7	5.6	7.7	5.7	8.6
necessary	MLP	0.4	1.1	0.3	1.0	0.6	1.5	0.6	1.1
	XGB	0.4	1.1	0.3	1.0	0.5	1.3	0.5	1.2
	RF	0.4	1.3	0.3	1.1	0.7	1.7	0.7	1.4

In both tasks, the length of necessary explanations is markedly lower than the length of sufficient explanations, highlighting that for many relations there are no necessary shared semantic aspects.

## EXAMPLES OF EXPLANATIONS FOR PPI PREDICTION


True Negative (-/-)

Proline-rich 5-like – Guanine nucleotide-binding 3-like

w/o 'regulation of nucleobase-containing compound metabolic process'

only 'protein binding'  
only 'regulation of nucleobase-containing compound metabolic process'


only sufficient  
w/o necessary  
global



## EXAMPLES OF EXPLANATIONS FOR PPI PREDICTION

False Negative (+/-)

Pulmonary surfactant-associated protein B – Granulocyte-macrophage colony-stimulating factor receptor subunit  $\alpha$



The aspects they share are very general, because both proteins are poorly described under the GO.