

Benchmark datasets for biomedical knowledge graphs with negative statements

Rita T. Sousa*, Sara Silva and Catia Pesquita

LASIGE, Faculdade de Ciências da Universidade de Lisboa

Abstract

Knowledge graphs represent facts about real-world entities. Most of these facts are defined as positive statements. The negative statements are scarce but highly relevant under the open-world assumption. Furthermore, they have been demonstrated to improve the performance of several applications, namely in the biomedical domain. However, no benchmark dataset supports the evaluation of the methods that consider these negative statements.

We present a collection of datasets for three relation prediction tasks - protein-protein interaction prediction, gene-disease association prediction and disease prediction - that aim at circumventing the difficulties in building benchmarks for knowledge graphs with negative statements. These datasets include data from two successful biomedical ontologies, Gene Ontology and Human Phenotype Ontology, enriched with negative statements.

We also generate knowledge graph embeddings for each dataset with two popular path-based methods and evaluate the performance in each task. The results show that the negative statements can improve the performance of knowledge graph embeddings.

Keywords

Biomedical Knowledge Graphs, Biomedical Ontologies, Gene Ontology, Human Phenotype Ontology, Negative Statements, Protein-Protein Interaction Prediction, Gene-Disease Association Prediction, Disease Prediction

1. Introduction


Knowledge Graphs (KGs) have been used to represent knowledge about real-world entities and their relationships. Most KGs use ontologies as a backbone to describe entities through ontology-based annotation, which associates an entity with a class. These annotations are commonly represented as positive statements establishing that an ontology class describes an entity. For example, in the biomedical domain, positive statements express that a protein $P1$ performs *intracellular_galactose_homeostasis* as defined in the Gene Ontology (GO) [1]. Negative statements are extremely rare but can be used to declare that a given protein $P2$ does not perform *glucose_homeostasis* (Figure 1).


SeWebMeDa-2023: 6th International Workshop on Semantic Web solutions for large-scale biomedical data analytics, May 29, 2023, Hersonissos, Greece

*Corresponding author.

✉ risousa@ciencias.ulisboa.pt/ (R. T. Sousa); sgsilva@ciencias.ulisboa.pt/ (S. Silva); clpesquita@ciencias.ulisboa.pt/ (C. Pesquita)

ORCID 0000-0002-7241-8970 (R. T. Sousa); 0000-0001-8223-4799 (S. Silva); 0000-0002-1847-9393 (C. Pesquita)

 © 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

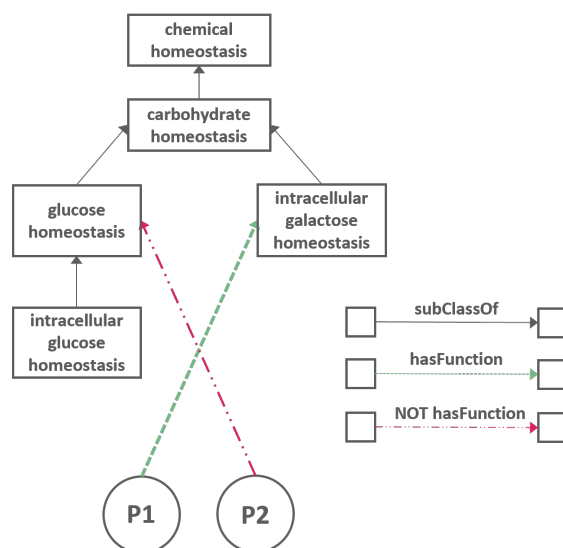


Figure 1: A GO KG subgraph with positive and negative statements describing two proteins.

The lack of negative statements is a significant issue because KGs operate under the open-world assumption. Therefore, this lack of information can lead to confusion regarding whether the absence of a positive statement is due to a lack of knowledge or the actual absence of the relationship. Moreover, the importance of negative statements to produce more accurate representations of entities in a KG [2, 3] and improving performance in different applications [4, 5] is increasingly recognized in the biomedical domain.

While there have been attempts to enhance current KGs with interesting negative statements, to the best of our knowledge, no benchmark datasets have been established to evaluate learning tasks over those KGs. With this in mind, we enrich existing biomedical KGs with negative statements and propose a collection of datasets for different biomedical tasks of relation prediction. The biomedical domain was selected because biomedical KGs are usually back-boned by biomedical ontologies that can express negation. Additionally, negative statements have been considered relevant for different biomedical applications [6]. Our datasets are grouped according to the task: protein-protein interaction (PPI) prediction, gene-disease association (GDA) prediction and disease prediction. Regarding the KGs, we enrich two successful biomedical ontologies: GO which covers distinct semantic aspects of gene products' function, and Human Phenotype Ontology (HP) which describes the universe of concepts related to phenotypic abnormalities found in human hereditary diseases.

2. Related Work

Several approaches to enriching existing KGs with interesting negative statements have been proposed. Arnaout *et al.* [7] proposed a method to enrich Wikidata by including interesting negative statements, which led to improvements in tasks involving entity summarization and decision-making.

Table 1

Statistics for each ontology regarding classes, nodes, edges.

	GO	HP
Classes	50918	17060
Literals and blank nodes	532373	442246
Edges	1425102	1082859

In the biomedical domain, several approaches tackle the lack of negative statements in biomedical ontologies, such as GO. The number of functions that a protein does not have is larger than the number of functions it has. Therefore, the number of negative statements describing proteins in the GO should be several orders of magnitude greater than the number of positive statements. Youngs *et al.* [8] designed two algorithms to predict negative statements for GO and populate the NoGo database, one based on empirical conditional probability and the other on topic modeling applied to genes and annotation. Fu *et al.* [4] introduced NegGOA, a new method to enrich the GO with relevant negative statements indicating that a protein does not perform a given function. This method exploits the GO by using hierarchical semantic similarity between GO terms. The enriched GO was used for protein function prediction. Later, Vesztröcy *et al.* [5] presented a benchmark based on a balanced test set of positive and negative statements. The negative statements are generated from expert-curated annotations of protein families on phylogenetic trees. The results of this work demonstrated that negative statements improve protein function prediction. Regarding the HP, although the importance of negative statements in gene-phenotype prediction is recognized, the enrichment with negative statements has yet to be investigated [2].

3. Building the Datasets

We present a collection of datasets that work over two enriched KGs for three relation prediction tasks: PPI prediction, GDA prediction, and disease prediction. Each benchmark dataset comprises several pairs of biomedical entities (or instances) that can be of the same type (protein-protein) or distinct types (gene-disease and disease-patient) with the respective label (1 for the positive pairs and zero for the negative pairs). Tables 1 and 2 show the KGs' and datasets' statistics for each task. Since for GDA prediction and disease prediction, the target relation happens between two types of instances (genes and diseases for GDA prediction and diseases and patients for disease prediction), the instance numbers in Table 2 appear separately. Moreover, in the case of PPI prediction, we exclusively employ the GO KG that has been subjected to a negative statement enrichment approach. However, when it comes to GDA prediction and disease prediction, we rely on the HP KG, which lacks a negative statement enrichment approach, resulting in a significant imbalance between the number of positive and negative statements.

To build these datasets, we adopt three main steps. The first one consists of enriching the KGs. The KG is constructed using the owlready2 package¹, which parses the ontology file

¹<https://owlready2.readthedocs.io/en/v0.37/>

Table 2

Statistics for each task's dataset regarding the number of instances, pairs, positive and negative statements.

	PPI prediction	GDA prediction	Disease Prediction
Instances	440	174 + 107	1033 + 660
Positive Pairs	1024	107	660
Negative Pairs	1024	107	681120
Positive statements	7364	14828	38130
Negative statements	8579	9191	179

```

<owl:NamedIndividual rdf:about="http://purl.obolibrary.org/obo/GO_0048268">
  <rdf:type rdf:resource="http://purl.obolibrary.org/obo/GO_0048268"/>
</owl:NamedIndividual>
<rdf:Description>
  <rdf:type rdf:resource="http://www.w3.org/2002/07/owl#NegativePropertyAssertion"/>
  <owl:sourceIndividual rdf:resource="http://purl.obolibrary.org/obo/GO_0048268"/>
  <owl:assertionProperty rdf:resource="http://purl.obolibrary.org/obo#has_function"/>
  <owl:targetIndividual rdf:resource="https://www.uniprot.org/uniprotkb/Q9BY11"/>
</rdf:Description>

```

Figure 2: Example of how the negative statements are defined in the OWL file.

in OWL format and processes the annotation file. The annotation file contains positive and negative statements used to describe entities. We use the guidelines established by the W3C² to define the negative statements as negative object property assertions³. To do so, we use metamodeling and represent each ontology class as a class and an individual. This situation translates into using the same IRI. Then, we use a negative object property assertion to state that the individual representing a biomedical entity is not connected by the object property expression to the individual representing an ontology class, as depicted in Figure 2. The second step consists of extracting pairs of entities from bioinformatic databases. The third step involves selecting the pairs containing KG entities that are well described with positive and negative statements.

The following subsections describe in more detail the KGs as the characteristics of each task.

3.1. Biomedical Knowledge Graphs

Two KGs back-boned by biomedical ontologies are used: the GO KG and the HP KG. Table 1 shows the statistics for each ontology.

The GO is used to describe gene products (proteins or genes) according to the molecular functions they perform, the biological processes they are involved in, and the cellular components where they act. The GO KG is built by integrating three sources: the GO⁴ itself, the GO

²<https://www.w3.org/TR/owl2-mapping-to-rdf/>

³https://www.w3.org/TR/owl2-syntax/#Negative_Object_Property_Assertions

⁴The GO was downloaded on September 2021. It is available at <http://release.geneontology.org/2021-09-01/ontology/index.html>

Annotation data⁵ [9], and negative GO associations produced in [5]⁶.

A GO annotation links a specific gene product with a particular GO class. The majority of GO annotation data corresponds to positive statements. However, the GO annotation has the qualifier 'NOT' for a few cases, meaning that a gene product has been proven not to carry out a specific function. The annotations that possess this qualifier were added as negative statements. In addition to these negative statements, the GO KG was also enriched with negative statements derived from expert-curated annotations of protein families on phylogenetic trees. The idea is that, if no evidence exists to suggest otherwise, gene function is maintained over time through evolution. Therefore, after expert curators have annotated ancestral states in gene phylogenies with GO classes, they check if the annotations are propagated down the phylogeny. When there is evidence that the function is absent in a specific sub-tree, a negative statement is added to that protein. These enriched negative statements were filtered so there were no contradictions with the GO annotation data.

HP characterizes phenotypic abnormalities discovered in human hereditary diseases according to five semantic aspects: phenotypic abnormalities, mode of inheritance, clinical course, clinical modifier and frequency. HP annotations can link diseases, patients or genes to HP classes via positive and negative statements. The construction of HP KG⁷ is similar to that of the GO KG. A negative annotation from HP that includes 'NOT' indicates that a disease does not cause that phenotype, so they are included as negative statements.

3.2. Protein-Protein Interaction Prediction Dataset

Predicting PPIs is a fundamental task in molecular biology for understanding biological systems. Given the high cost of experimentally determining PPI, many computational approaches for PPI prediction based on available functional information described by the GO [1] have been proposed to find protein pairs likely to interact and thus provide a selection of good candidates for experimental analysis. Therefore, the GO KG is used to describe the proteins of the dataset.

The positive examples are extracted from the STRING [10] database. Our selection of protein pairs was based on the following criteria: (i) interactions between proteins had to be curated or experimentally determined rather than computationally determined; (ii) interactions needed to have a confidence score above 0.950 to ensure high confidence; (iii) each protein must have at least one positive statement for a GO class and one negative statement for another GO class. The negative examples are generated by random negative sampling over the set of proteins of the positive examples.

3.3. Gene-Disease Association Prediction Dataset

Knowing which genes are associated with a specific disease is crucial to understanding the disease mechanisms and recognising potential biomarkers or therapeutic targets. However, once again, validating these associations in the wet lab is expensive and time-consuming. This has

⁵The GO positive annotations were downloaded on January 2021. It is available at <http://release.geneontology.org/2021-01-01/annotations/index.html>.

⁶The negative annotations were downloaded from https://lab.dessimoz.org/20_not

⁷The HP was downloaded on October 2022, while the HP annotations were downloaded on November 2021. A link to these versions is no longer available.

prompted the evolution of computational methods to identify the most promising associations to be further validated.

The two KGs are used for the GDA prediction task dataset. GO KG describes the genes, and HP KG describes the diseases. The target relations to predict are extracted from DisGeNET [11]. Adapting the approach described in [12], we considered the following criteria to select gene-disease pairs: (i) each gene must have at least one positive statement for a GO class and one negative statement for another GO class; (ii) each disease must have at least one positive statement for an HP class and one negative statement for an HP class. We sampled random negative examples of the same genes and diseases to create a balanced dataset.

3.4. Disease Prediction Datasets

Since human diseases are a complex phenomenon, disease prediction is an essential but still complicated task that must be executed accurately and efficiently. Therefore, using computational methods to help physicians prioritize diseases is highly advantageous.

The dataset to predict if a synthetic patient has been diagnosed with a specific disease is generated by adapting the methodology proposed in [13]. Thirty-three mendelian diseases for which they knew the penetrance of each phenotype are selected. Penetrance indicates the likelihood that a patient suffering from a specific disease will exhibit a particular phenotype. For each of these 33 diseases, 20 synthetic patients diagnosed with that disease are created. The patients' positive annotation is determined by the disease's penetrance and the patient's gender. The gender is defined randomly with an equal likelihood for both genders. For example, the 'Aarskog-Scott syndrome' is annotated with the phenotype 'Ptosis' with a penetrance of 0.5061, meaning that approximately half of the synthetic patients diagnosed with that disease will have a positive statement for this phenotype. The negation of phenotypes does not have a penetrance associated, so synthetic patients inherit the negative phenotypes related to the disease. For example, since the disease 'Aarskog-Scott syndrome' is annotated with 'NOT Decreased Fertility', each patient will have a negative statement for this phenotype. Furthermore, 1000 diseases were randomly chosen to add complexity to the task. These diseases are annotated with positive and negative statements.

Random annotations can also be added to patients to emulate a more realistic situation where a patient is associated with phenotypes unrelated to the patient's disease. In addition to the disease prediction dataset, we present three versions with random annotations. The number of random annotations is defined by a percentage Noi ($\text{Noi}=[0, 0.1, 0.2, 0.4]$) concerning a given patient's total number of annotations. For example, if $\text{Noi}=0.5$, half of the full annotations of a given patient are added. Table 3 shows the number of positive and negative statements for each noise version.

4. Validation of the Datasets

KG embedding methods [14] have been successfully employed in several biomedical applications [14]. Since these methods map KGs into low-dimensional spaces, they have emerged as a popular way to generate features for machine learning tasks. Therefore, we use two KG embedding methods to evaluate our datasets - RDF2Vec [15] and OWL2Vec* [16]. RDF2Vec

Table 3

Statistics for each noise version regarding the number of positive and negative statements.

Noise	Positive Statements	Negative Statements
0.1	40592	216
0.2	37195	214
0.4	38242	192

Table 4

Median precision, recall and weighted average F-measure (Pr, Re, and F1) for PPI, GDA, and disease prediction using only positive statements (Pos) or positive and negative statements (Pos+Neg).

Method	Stat.	PPI			GDA			Disease Prediction		
		Pr	Re	F1	Pr	Re	F1	Pr	Re	F1
RDF2Vec	Pos	0.831	0.826	0.828	0.623	0.625	0.615	0.994	0.742	0.850
	Pos+Neg	0.847	0.844	0.845	0.654	0.600	0.645	1.000	0.771	0.870
OWL2Vec*	Pos	0.833	0.806	0.823	0.652	0.656	0.646	0.975	0.584	0.730
	Pos+Neg	0.860	0.812	0.840	0.625	0.661	0.630	0.980	0.563	0.713

is a path-based method that generates random walks in the KG that constitutes the corpus of word sequences given as input to a neural language model. OWL2Vec* was designed to learn ontology embeddings and it also employs direct walks on the graph to learn graph structure. These embedding methods generate representations of the biomedical entities that are combined using the binary Hadamard operator to represent the pair.

The pair representations are then fed into a Random Forest algorithm for training using Monte Carlo cross-validation (MCCV) [17]. MCCV is a variation of traditional k -fold cross-validation in which the data is divided into training and testing sets (with β being the proportion of the dataset to include in the test split) M times. Our experiments use MCCV with $M = 30$ and $\beta = 0.3$ for PPI and GDA prediction. Given the large number of pairs for disease prediction, we use MCCV with $M = 5$ and $\beta = 0.3$.

Each embedding method is run with two different KGs, one with only positive statements and the other with both positive and negative statements. Table 4 reports each task’s median of recall, precision and weighted average F-measure.

Figure 3 compares the impact of using only positive statements versus both positive and negative statements on our datasets. The bars represent the difference in performance for precision, recall and weighted average F-measure, with upward bars indicating improved performance with both positive and negative statements and downward bars indicating decreased performance.

The experiments show that the added information given by negative statements generally improves the performance of RDF2Vec. However, for OWL2Vec*, the performance only improves for PPI prediction.

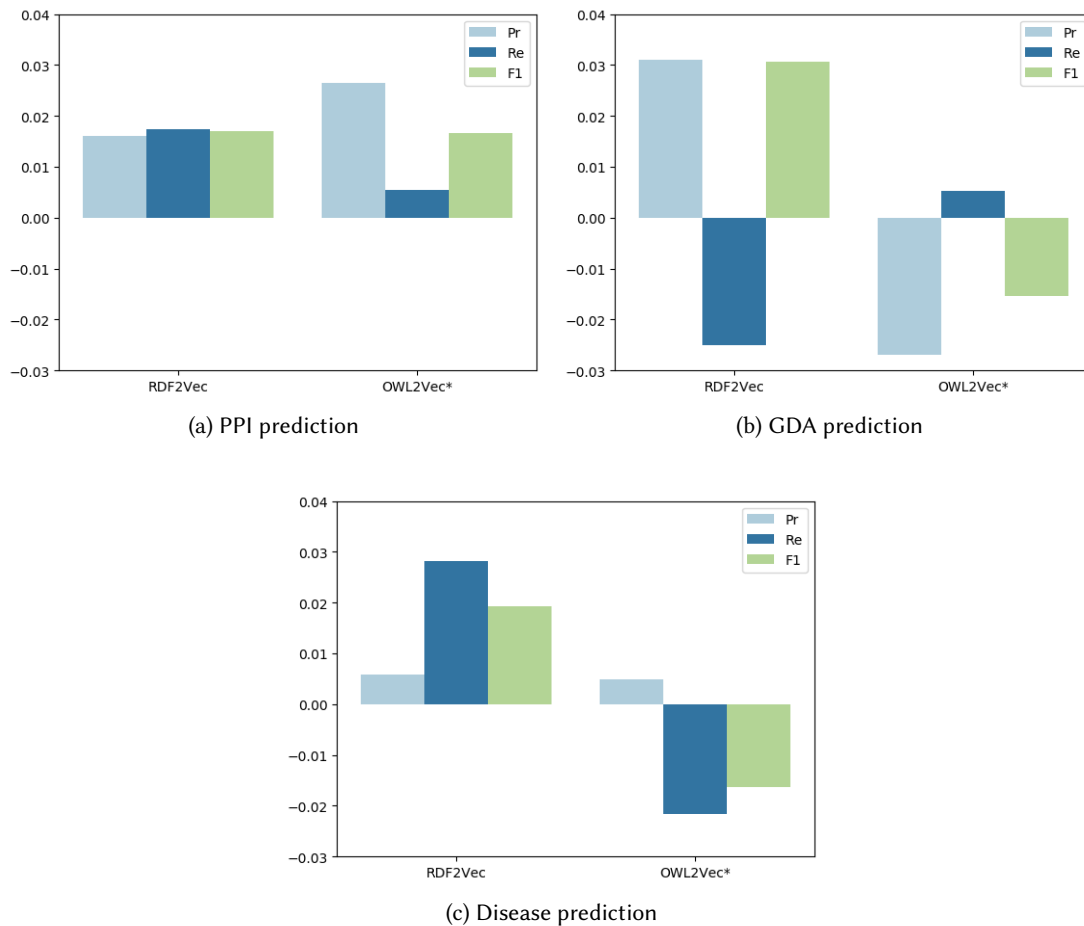


Figure 3: Barplots showing the differences in precision, recall, weighted average F-measure (Pr, Re, and F1) between using positive and negative statements or only positive statements.

5. Using the Benchmark

All datasets are available on Zenodo⁸ under a CC BY 4.0 license. For each dataset, we provide access to two types of files: (1) one TSV file containing pairs of entities and information about whether a relationship exists between them or not; (2) OWL files containing the KG used to describe the biomedical entities that appear in the TSV file. Together, these files can be used to perform relation prediction tasks since the TSV file provides the specific entities and relations that need to be predicted, while the OWL file provides the necessary background knowledge for generating the features.

⁸<https://doi.org/10.5281/zenodo.7709195>

6. Conclusions

Benchmark datasets are essential for evaluating and comparing the performance of different approaches that work over KGs. This paper presents a collection of datasets for three relation prediction tasks in the biomedical domain: PPI prediction, GDA prediction, and disease prediction. The biomedical domain is chosen since it is already demonstrated that the inadequacy of approaches to take into consideration negative statements is a limitation for several biomedical applications. However, although the datasets are domain-specific, they can be used to evaluate approaches outside the biomedical domain.

The datasets are validated using two popular KG embedding methods to generate features that are then given as input for a classifier. The results highlight the importance of incorporating negative statements into KGs to create more accurate representations of KG entities.

Acknowledgments

C. P., S. S., R. T. S. are 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, which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101017453, and by HfPT: Health from Portugal under the Portuguese Plano de Recuperação e Resiliência. The authors thank Lina Aveiro for the preliminary results of this work.

References

- [1] GO Consortium, The Gene Ontology Resource: 20 years and still GOing strong, *Nucleic Acids Research* 47 (2018) D330–D338.
- [2] L. Liu, S. Zhu, Computational methods for prediction of human protein-phenotype associations: A review, *Phenomics* 1 (2021) 171–185.
- [3] P. Gaudet, C. Dessimoz, Gene Ontology: pitfalls, biases, and remedies, in: *The Gene Ontology Handbook*, Humana Press, New York, NY, 2017, pp. 189–205.
- [4] G. Fu, J. Wang, B. Yang, G. Yu, NegGOA: negative GO annotations selection using ontology structure, *Bioinformatics* 32 (2016) 2996–3004.
- [5] A. Warwick Vesztrocy, C. Dessimoz, Benchmarking Gene Ontology function predictions using negative annotations, *Bioinformatics* 36 (2020) i210–i218.
- [6] M. Kulmanov, F. Z. Smaili, X. Gao, R. Hoehndorf, Semantic similarity and machine learning with ontologies, *Briefings in Bioinformatics* 22 (2021) bbaa199.
- [7] H. Arnaout, S. Razniewski, G. Weikum, J. Z. Pan, Negative statements considered useful, *Journal of Web Semantics* 71 (2021) 100661. Publisher: Elsevier.
- [8] N. Youngs, D. Penfold-Brown, R. Bonneau, D. Shasha, Negative example selection for protein function prediction: The NoGO database, *PLOS Computational Biology* 10 (2014) 1–12.
- [9] GO Consortium, The Gene Ontology resource: enriching a GOld mine, *Nucleic Acids Research* 49 (2021) D325–D334.

- [10] D. Szklarczyk, A. L. Gable, K. C. Nastou, D. Lyon, R. Kirsch, S. Pyysalo, N. T. Doncheva, M. Legeay, T. Fang, P. Bork, L. J. Jensen, C. von Mering, The STRING database in 2021: customizable protein–protein networks, and functional characterization of user-uploaded gene/measurement sets, *Nucleic Acids Research* 49 (2020) D605–D612.
- [11] J. Piñero, J. M. Ramírez-Anguita, J. Saüch-Pitarch, F. Ronzano, E. Centeno, F. Sanz, L. I. Furlong, The DisGeNET knowledge platform for disease genomics: 2019 update, *Nucleic Acids Research* 48 (2019) D845–D855.
- [12] S. Nunes, R. T. Sousa, C. Pesquita, Predicting gene–disease associations with knowledge graph embeddings over multiple ontologies, in: *ISMB Annual Meeting - Bio-Ontologies*, 2021.
- [13] A. J. Masino, E. T. Dechene, M. C. Dulik, A. Wilkens, N. B. Spinner, I. D. Krantz, J. W. Pennington, P. N. Robinson, P. S. White, Clinical phenotype-based gene prioritization: an initial study using semantic similarity and the human phenotype ontology, *BMC bioinformatics* 15 (2014) 1–11.
- [14] Q. Wang, Z. Mao, B. Wang, L. Guo, Knowledge graph embedding: A survey of approaches and applications, *IEEE Transactions on Knowledge and Data Engineering* 29 (2017) 2724–2743.
- [15] P. Ristoski, H. Paulheim, RDF2Vec: RDF graph embeddings for data mining, in: *International Semantic Web Conference*, Springer, 2016, pp. 498–514.
- [16] J. Chen, P. Hu, E. Jimenez-Ruiz, O. M. Holter, D. Antonyrajah, I. Horrocks, OWL2Vec*: Embedding of OWL ontologies, *Machine Learning* (2021) 1–33.
- [17] Q.-S. Xu, Y.-Z. Liang, Monte Carlo cross validation, *Chemometrics and Intelligent Laboratory Systems* 56 (2001) 1–11.