

RESEARCH

Open Access



BioBLP: a modular framework for learning on multimodal biomedical knowledge graphs

Daniel Daza^{1,2,4*}, Dimitrios Alivanistos^{1,4*}, Payal Mitra³, Thom Pijenburg³, Michael Cochez^{1,4} and Paul Groth^{2,4}

Abstract

Background Knowledge graphs (KGs) are an important tool for representing complex relationships between entities in the biomedical domain. Several methods have been proposed for learning embeddings that can be used to predict new links in such graphs. Some methods ignore valuable attribute data associated with entities in biomedical KGs, such as protein sequences, or molecular graphs. Other works incorporate such data, but assume that entities can be represented with the same data modality. This is not always the case for biomedical KGs, where entities exhibit heterogeneous modalities that are central to their representation in the subject domain.

Objective We aim to understand how to incorporate multimodal data into biomedical KG embeddings, and analyze the resulting performance in comparison with traditional methods. We propose a modular framework for learning embeddings in KGs with entity attributes, that allows encoding attribute data of different modalities while also supporting entities with missing attributes. We additionally propose an efficient pretraining strategy for reducing the required training runtime. We train models using a biomedical KG containing approximately 2 million triples, and evaluate the performance of the resulting entity embeddings on the tasks of link prediction, and drug-protein interaction prediction, comparing against methods that do not take attribute data into account.

Results In the standard link prediction evaluation, the proposed method results in competitive, yet lower performance than baselines that do not use attribute data. When evaluated in the task of drug-protein interaction prediction, the method compares favorably with the baselines. Further analyses show that incorporating attribute data does outperform baselines over entities below a certain node degree, comprising approximately 75% of the diseases in the graph. We also observe that optimizing attribute encoders is a challenging task that increases optimization costs. Our proposed pretraining strategy yields significantly higher performance while reducing the required training runtime.

Conclusion BioBLP allows to investigate different ways of incorporating multimodal biomedical data for learning representations in KGs. With a particular implementation, we find that incorporating attribute data does not consistently outperform baselines, but improvements are obtained on a comparatively large subset of entities below a specific node-degree. Our results indicate a potential for improved performance in scientific discovery tasks where understudied areas of the KG would benefit from link prediction methods.

Keywords Multimodal knowledge graph, Graph embeddings, Biomedical knowledge

*Correspondence:

Daniel Daza

d.dazacruz@vu.nl

Dimitrios Alivanistos

d.alivanistos@vu.nl

Full list of author information is available at the end of the article



© The Author(s) 2023. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>. The Creative Commons Public Domain Dedication waiver (<http://creativecommons.org/publicdomain/zero/1.0/>) applies to the data made available in this article, unless otherwise stated in a credit line to the data.

Background

The domain of Life Sciences (LS) relies on large scale and diverse data; ranging from genomics and proteomics data, to electronic health care records, and adverse events reports [1–3].

A critical challenge to effectively using this data and inferring complex relationships across different sources is that it is often not interoperable [4] and it cannot be easily integrated, hindering the ability to query and analyze the available knowledge [5]. To address this problem, semantic technologies have been employed to create large scale LS Knowledge Graphs [5–9]. Knowledge Graphs (KG) are graph structured knowledge bases of entities and their relations [10], enabling, for example, the study of the relationship between chemical compounds and secondary pharmacology [11] as well as drug repurposing [9]. Examples of KGs built for the LS include Bio2RDF [6], Open PHACTS [7], and Hetionet [9], which represent extensive repositories of various relationships between entities in the biomedical domain, as well as descriptive attributes about entities such as molecular structures, amino acid sequences, and disease descriptions. This information can be complemented with additional data from resources such as DrugBank [12], UniProt [13], and MeSH [14].

The use of entity attributes brings about the possibility of incorporating a wealth of multimodal domain-specific data into the already extensive information present in the graph structure of the KG. Crucially, attribute data can have a complementary relationship with the information in the graph: while two proteins might have the same relations with given molecules, their amino acid sequences might be different, as we illustrate in Fig. 1. Having access to attribute data can thus be fundamental for distinguishing entities that are identical according to the topology of the KG.

We are interested in knowledge graphs where entities can be associated with attributes of specific modalities (e.g., amino acid sequences for proteins), that encode important entity attributes. Formally, we define such a KG as a tuple $G = (V, R, E, D, d)$, where V is a set of nodes, R a set of relations, E a set of triples (or *edges*) of the form (h, r, t) with $h, t \in V$ and $r \in R$. We refer to h as the head, r as the relation, and t as the tail of the triple. D is a dataset containing attribute data of entities, and d a partial function $d : V_d \rightarrow D$ where $V_d \subseteq V$ is the subset of entities with attribute data, and $d(v_i)$ retrieves the attribute data of entity v_i . We define d as a partial function to formally support KGs for which attribute data is available for only a subset of the entities, which is a common scenario when dealing with large and heterogeneous graphs.

The information contained in KGs is often incomplete because certain facts are unknown or not considered important enough to make explicit. In other cases, changes in the domain could require the addition of new entities, new edges, or updates to existing ones. This affects their ability to support the derivation of insights and scientific discovery. Instead of retrieving existing relations between entities in the graph, a scientist might be interested in relations that are not yet present, but are nonetheless *likely* [15]. These inferred relations could lead to the recommendation of a new hypothesis to be evaluated, such as an untested interaction between a protein and a molecule.

To tackle this scenario, different approaches have been proposed in machine learning literature that seek to exploit the patterns in the graph to train *link prediction models*, which generalize to edges not observed during training. In this work, we are concerned with the problem of link prediction over biomedical KGs with entity attributes. Various challenges arise when optimizing link prediction models from both graph structure and entity attributes, especially in the biomedical domain.

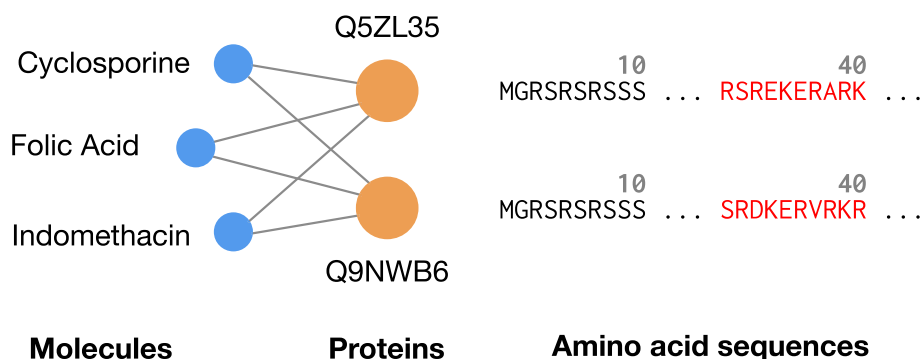


Fig. 1 Two proteins (labeled with UniProt identifiers) from the BioKG dataset, which interact with the same molecules. When viewed from the perspective of the edges in the knowledge graph, the proteins are indistinguishable from each other. Their respective amino acid sequences on the right constitute a valuable signal that reflects their differences

The multimodal nature of entity attribute data imposes a computational burden on models, further complicated by the fact that the coverage of attribute data is rarely perfect across the entire graph. We propose a modular framework for learning knowledge graph embeddings that incorporates entity domain attributes in a link prediction setting that is subjected to these challenges.

Link prediction with knowledge graph embeddings

Knowledge graph embeddings have become an increasingly popular tool for link prediction [16, 17]. They consist of a mapping from the entities and relations in the graph, to vectors in a space where functions are used to model the likelihood of a triple.

More formally, KG embeddings can be defined in terms of an embedding function $e : E \cup R \rightarrow \mathcal{X}$ that maps an entity or relation in the graph to an element of an embedding space \mathcal{X} , plus a scoring function $f : \mathcal{X}^3 \rightarrow \mathbb{R}$ that given the embedding of the entities and relation in a triple, computes a score that indicates the likelihood of a triple. Denoting the embedding of entities and relations in boldface (e.g. $e(h) = \mathbf{h}$), the embeddings are optimized so that the score $f(\mathbf{h}, \mathbf{r}, \mathbf{t})$ is high for a triple (h, r, t) that is part of the KG, and low for a triple that is not.

A large part of the research carried out in the field of KG embeddings is concerned with designing suitable spaces \mathcal{X} alongside embedding and scoring functions that are scalable and expressive enough to learn patterns that might occur in a KG. In the simplest of cases, the embedding function is a lookup operation that for each entity and relation, returns an element of \mathcal{X} . We call these methods *look-up table embeddings*, where the table corresponds to a matrix of embeddings and a lookup or index that can be used to retrieve a specific row from it. A well known example of such a method is TransE [18], which maps entities and relations to vectors in a Euclidean space \mathbb{R}^n of some dimension n . For a triple (h, r, t) , TransE employs a translational model such that the sum of embeddings $\mathbf{h} + \mathbf{r}$ is close to \mathbf{t} . This results in a scoring function that computes how close the vectors are:

$$f_{\text{TransE}}(\mathbf{h}, \mathbf{r}, \mathbf{t}) = -\|\mathbf{h} + \mathbf{r} - \mathbf{t}\|. \quad (1)$$

Other examples of lookup table embeddings are ComplEx [19] and RotatE [20], which instead use vectors defined over a complex-valued space \mathbb{C}^n , and scoring functions that perform multiplications or rotations over such vectors. While these models have shown promising results in the task of predicting missing links between entities (also known as the *link prediction* task), they suffer from two fundamental limitations.

First, they are optimized using only the structure of the graph, discarding other informative signals. In particular,

these representations ignore important information such as entity attributes that, as we outlined before, can contain critical information for differentiating between two entities in the graph. Second, the mapping from the entities and relations to some embedding in the lookup table has to be fixed before running the training algorithm. If new entities are added to the KG, the embeddings for these entities are not defined, which results in the inability to predict any links involving the new entities. This limitation is at odds with the dynamic nature of KGs, especially in the biomedical domain.

Knowledge graph embedding methods that can compute predictions for entities not seen during training are known as *inductive* link prediction models [21–24]. An effective method to achieve this is to rely on entity attributes, which we will discuss next.

Incorporating attributes in knowledge graph embeddings

Instead of directly mapping entities in a KG to an embedding, some methods have proposed embedding functions that take as input the attributes of an entity [21, 23, 25–31]. The embedding function can then be thought of as an *encoder*, that can be designed through the use of neural network architectures to fit to specific modalities of the attribute data. The task of the encoder is to process this data and output the embedding of the entity, which can in turn be passed to a particular scoring function. During training, the encoder is optimized to map entity attributes to embeddings that are helpful for predicting links between entities. Once training is finished, the encoder is able to produce embeddings for entities not seen during training as long as they are paired with their respective attribute data, thereby achieving the goal of inductive link prediction.

Recent methods that incorporate attributes assume that all entities in the graph are paired with a textual description, which is turned into a sequence of word embeddings and passed to a neural network that processes and aggregates the embeddings into a single vector [23, 32–34]. Moreover, recent methods rely on pretrained language models such as BERT [35] as the backbone of the attribute encoder. This allows the KG embeddings to benefit from the expressive architecture of such language models, and the encoded knowledge they contain through the vast amount of data used to pretrain them.

While previous methods allow encoding entity attributes into embeddings, they work on two fundamental assumptions: i) that all the attributes of entities share the same modality (e.g. text), and ii) that such information is available for all entities in the graph. Biomedical KGs, especially those composed of data from different original sources are, on the contrary, characterized by heterogeneous entity types, each of which is associated with data of different modalities. For example, disease entities can be described

using text, molecules can be represented as small graphs of atoms and bonds, and proteins as sequences of amino acids. It is therefore unclear how to directly transfer the results of encoder-based KG embeddings designed for text attributes, to the case of biomedical KGs where a textual representation might be natural only for some types of entities. Furthermore, such methods do not provide a mechanism for dealing with entities with missing attribute data.

Knowledge graph embedding models in biology

KG embeddings can be used to tackle many LS problems, such as drug repurposing and target identification in the drug development process. Previous works have shown that this can be achieved by collecting a sufficiently large knowledge graph that links entities such as proteins, diseases, genes, and drugs, and then training a model for various prediction tasks [36–42]. In particular, some of these works transfer the findings from the literature of KG embeddings—especially lookup table embeddings—to the biomedical domain, finding that the embeddings are effective at solving problems such as Drug-Target Interaction (DTI) prediction [40, 41, 43].

While these results are promising, prior work has also highlighted challenges for using KG embeddings in practice [36, 38]. In particular, KG embeddings perform better on well-studied biomedical entities, because more detailed information is available about them that can be incorporated in the graph. Conversely, this can cause low link prediction performance on *low degree* entities; i.e. entities with relatively few connections to other entities in the graph. Moreover, by using lookup table embeddings, these methods are still limited to the structural information contained in the graph, and do not allow for predictions involving new entities [40].

Recent work has investigated the effect of incorporating entity attributes into embeddings for various applications in the biomedical domain [44–49]. The proposed methods are specialized to certain data modalities, such as textual descriptions [44, 45, 49], or SMILES molecular representations [46–48]. These methods are designed for biomedical KGs where the *data modality is the same for all entities* in the graph, and they also assume that attribute data is available for all entities. This requirement makes existing unimodal methods unsuitable for direct comparison with BioBLP framework on KGs such as BioKG which by definition comprise of heterogeneous entities that do not share attributes of the same modality. For instance, the same attribute encoder for protein sequences could not be used to encode disease entities since disease entities do not share the same attribute. Furthermore, some works have found that incorporating attribute data can outperform baseline methods based on the structure of the graph only [45–48]. However, these works do not

disclose whether or not hyperparameter optimization of baselines was carried out [45–47], or indicate that it was omitted [48]. We argue that adequate hyperparameter optimization is crucial for understanding the effect of incorporating attributes in comparison with methods based on the graph structure only, given that the latter can be quite sensitive to hyperparameters [50–52].

Research objective

Motivated by the limitations in the literature in relation to embedding multimodal knowledge graphs, we aim to answer the following research questions:

- **RQ1:** How can we incorporate multimodal attribute data into KGE models under a common learning framework?
- **RQ2:** Given this framework, what is the performance in comparison with baselines that do not incorporate attribute data on link prediction and drug-protein interaction benchmarks?
- **RQ3:** Under what conditions do we observe positive or negative effects of incorporating attribute data using our proposed framework?

By answering these questions, we seek to obtain an understanding of the practical implications of combining KGEs with attribute encoders in the context of biomedical data.

Methods

In the following sections, we describe the biomedical KG used in our research, the proposed framework for learning embeddings, and the evaluation protocols used to assess its performance.

The BioKG knowledge graph

BioKG is a Knowledge Graph built for the purpose of supporting relational learning in complex biological systems [38], with a focus on protein and molecules (interchangeably referred to as drugs). It was created in response to the problem of data integration and linking between various open biomedical knowledge bases, resulting in a compilation of curated data from 13 different biomedical databases, including DrugBank [12], UniProt [13], and MeSH [14].

The pipeline for creating BioKG employs a module to programmatically retrieve, clean and link the different entities and relations found in different databases, by reconciling their identifiers. This results in a biomedical KG containing over 2 million triples, where each entity has a uniform identifier linking it to publicly available databases. The coverage over multiple biomedical databases in BioKG presents a significant potential for training machine learning models on KGs with attribute data.

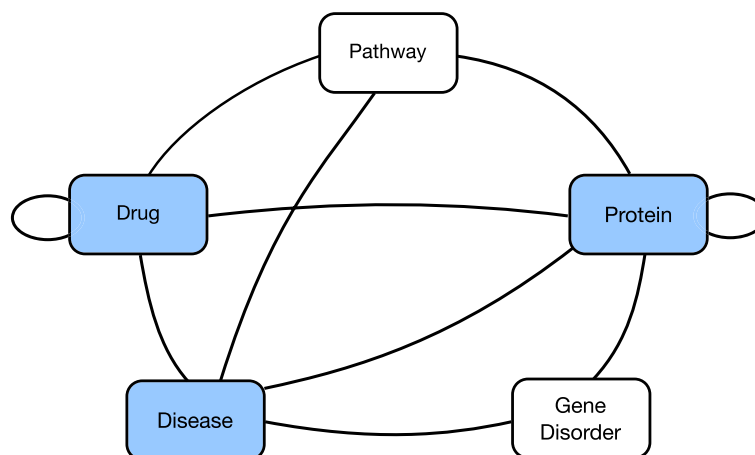


Fig. 2 Visualization of the entity types in the BioKG graph. Edges between nodes denote the existence of any relation type. Shaded nodes indicate entities for which we retrieve attribute data

Additionally, the availability of identifiers in BioKG facilitates the search for entity attributes of different modalities. These characteristics make BioKG an ideal choice for furthering research in the direction of machine learning on knowledge graphs with attribute data.

BioKG contains triples involving entities of 5 types: proteins, drugs, diseases, pathways, and gene disorders. We use the publicly available repository¹ to compile BioKG from the original sources, as available on May 2022. A visualization of the schema is shown in Fig. 2, and the statistics of the graph are listed in Table 1.

Following our aim of learning embeddings that incorporate attribute data of different modalities, while taking into account entities with missing attribute data, we retrieve attributes for proteins, drugs, and diseases. To do so, we take advantage of the uniform entity identifiers present in BioKG, that allow us to query publicly available API endpoints. In particular, we retrieve amino acid sequences from Uniprot², SMILES representations for drugs from DrugBank³, and textual descriptions of diseases from MeSH⁴. In a few cases where a textual MeSH scope note was unavailable for a disease, we use the text label for the disease as is, as this can contain useful information about the class, or anatomical categorization of a disease. We report the resulting coverage of attribute data for the entity types of our interest in Table 1. The result is a KG containing attribute data of three modalities for approximately 68% of the entities, while for the rest of the entities attribute data is missing, which are the desired characteristics of the KG for our research objective.

Table 1 Statistics of the BioKG graph used in our work. The Attribute Coverage refers to the number of entities that have associated attribute data, which we retrieve for proteins, molecules, and diseases

Statistic	Count	Attribute coverage
Edges	2,074,346	—
Relations	17	—
Entity types	5	—
Entities	106,337	73,131
Proteins	59,664	59,551
Drugs	8,808	7,766
Diseases	5,814	5,814
Other	32,501	0

Biomedical benchmarks

One aspect that separates BioKG from other related efforts [7], is the inclusion of benchmark datasets around domain-specific tasks such as Drug-Protein Interaction and Drug-Drug Interaction prediction. This is crucial, as it allows for experimentation and comparison to standard benchmark datasets and approaches. The benchmarks consist of pairs of entities for which a particular relationship holds, such as an interaction between drugs. We use them for comparing the predictive performance of KG embeddings obtained when using the graph structure alone, and when incorporating attribute data. In total, five benchmarks are provided with BioKG, all considering different drug-drug and drug-protein interactions.

We focus on the DPI-FDA benchmark, which consists of drug target protein interactions of FDA approved drugs compiled from KEGG and DrugBank databases. It is comprised of 18,928 drug-protein interactions of 2,277 drugs and 22,654 proteins. This benchmark is

¹ Available at <https://github.com/dsi-bdi/biokg>

² <https://www.uniprot.org/>

³ <https://go.drugbank.com/>

⁴ <https://www.nlm.nih.gov/mesh/>

Table 2 Statistics of the BioKG splits of triples used for training, validation, and testing, after removing triples overlapping with the benchmarks

	Total	Training	Validation	Test
BioKG	1,852,262	1,481,809	185,226	185,227

an extension of the DrugBank_FDA [53] and Yamani-shi09 [54] datasets which have 9,881 and 5,127 DPIs respectively. We refer to it as the DPI benchmark.

Data partitioning

In order to train and evaluate different KG embedding methods, we need to specify sets of triples for training, validation, and testing. Considering that the original BioKG dataset does not provide predefined splits, we generate our own. Since we aim to train multiple embedding methods on BioKG, and evaluate their performance over the benchmarks, we must guarantee that the KG and the benchmarks are *decoupled*, so that there is no data leakage from the benchmarks into the set of KG triples used for training.

Our strategy for partitioning the data is as follows:

1. We merge the five benchmarks provided in BioKG into a single collection of entity pairs.
2. For every pair of entities (x, y) in the collection, we remove all triples from BioKG in which x and y appear together, regardless of which one is at the head or tail of the triple.
3. Removing triples from BioKG in step 2 might lead to the removal of an entity entirely from the graph.
4. We split the resulting triples in BioKG into training, validation, and test sets, ensuring that all entities occur at least once in the training set of triples.

The last two steps are a requirement for lookup table embeddings. These methods require that all entities are observed in the training set of triples, so that we can obtain embeddings for evaluating i) link prediction performance on the validation and test triples, and ii) prediction accuracy on the DPI benchmark.

The proportions used for selecting training, validation, and test triples are 0.8, 0.1, and 0.1, respectively. We present statistics of the generated splits in Table 2. In the DPI benchmark, we end up with 18,678 pairs that we used to train classifiers and evaluate them via cross-validation.

To further research in this area, we make all the data we collected publicly available, comprising the attribute data for proteins, molecules, and diseases, and the splits generated for training, validation, and testing of machine learning models trained on BioKG⁵.

BioBLP

We now detail our new method, BioBLP, proposed as a solution to the problem of learning embeddings of entities in a KG with attributes. We design it as an extension of BLP [23] to the biomedical domain that incorporates the three data modalities available to us in the BioKG graph. BLP is a method for embedding entities in KGs where all entities have an associated textual description. At its core lies an encoder of textual descriptions based on BERT, a language model pretrained on large amounts of text [35]. In this work, we propose a more general formulation that considers multiple modalities as well as missing attribute data. By supporting KGs that contain entities *with* and *without* attributes, we can exploit the full extent of the information contained in the graph, rather than limiting the data available for training to one of the two cases.

BioBLP is a model for learning embeddings of entities and relations in a KG with attribute data of different modalities. It is comprised by four components: a set of attribute encoders, lookup table embeddings for entities with missing attributes and for relation types, a scoring function, and a loss function.

Attribute encoders Let $G = (V, R, E, D, d)$ be a KG with attribute data. BioBLP contains a set of k modality specific encoders $\{e_1, \dots, e_k\}$. Each of the e_i are functions of the attributes of an entity that output a fixed-length embedding. If we denote as $D_i \subset D$ the subset of attribute data of modality i (e.g. the subset of protein sequences), then formally each encoder is a mapping $e_i : D_i \rightarrow \mathcal{X}$, where \mathcal{X} is the embedding space shared by all entities. We design the encoders as neural networks with learnable parameters, whose architecture is tailored to the input modality.

Lookup table embeddings For entities that are missing attribute data, we employ a lookup table of embeddings. For each entity in V with missing attributes, we assign it an initial embedding by randomly sampling a vector from \mathcal{X} . We follow the same procedure for all relations in R . These embeddings are then optimized, together with the parameters of the attribute encoders in F . In practical terms, this means that when attributes are not available at all, BioBLP reduces to lookup table embedding methods such as TransE [18]. We implement the lookup table embeddings via matrices \mathbf{W}_e and \mathbf{W}_r for entities and relations, respectively, with each of the rows containing a vector in \mathcal{X} . The embedding function of an entity without attributes is an additional encoder e_0 such that $e_0(h)$ returns the row corresponding to entity h in \mathbf{W}_e . Similarly for relations, we define its embedding via a function e_r , the retrieves the embedding of a relation from the rows of \mathbf{W}_r .

Scoring function Once the attribute encoders and the lookup table embeddings have been defined, for a given triple (h, r, t) we can obtain the corresponding embeddings $\mathbf{h}, \mathbf{r}, \mathbf{t} \in \mathcal{X}$. These can be passed to any scoring function for

⁵ Available at <https://doi.org/10.5281/zenodo.8005711>.

triples $f : \mathcal{X}^3 \rightarrow \mathbb{R}$, such as the TransE scoring function defined in Eq. 1. The scoring function also determines the type of embedding space \mathcal{X} . For TransE, this corresponds to the space of real vectors \mathbb{R}^n while for ComplEx and RotatE it is the space of complex vectors \mathbb{C}^n .

Loss function The loss function is designed so that it can be minimized by assigning high scores to observed triples, and low scores to corrupted triples [51]. Corrupted triples are commonly obtained by replacing the head or tail entities in a triple (h, r, t) by an entity sampled at random from the KG. For such a triple, the loss function can thus be defined as a function $\mathcal{L}(f(\mathbf{h}, \mathbf{r}, \mathbf{t}), f(\tilde{\mathbf{h}}, \mathbf{r}, \tilde{\mathbf{t}}))$. An example is the margin ranking loss function:

$$\mathcal{L}(f(\mathbf{h}, \mathbf{r}, \mathbf{t}), f(\tilde{\mathbf{h}}, \mathbf{r}, \tilde{\mathbf{t}})) = \max(0, f(\tilde{\mathbf{h}}, \mathbf{r}, \tilde{\mathbf{t}}) - f(\mathbf{h}, \mathbf{r}, \mathbf{t}) + m), \quad (2)$$

which forces the score of observed triples to be higher than for corrupted triples by a margin of m or more.

We optimize the parameters of the attribute encoders and the lookup table embeddings via gradient descent on the loss function. If we denote by θ the collection of parameters in the entity encoders plus the lookup table embeddings, the training update rule during for a specific triple is the following:

$$\theta^{(t+1)} = \theta^{(t)} - \alpha \nabla_{\theta} \mathcal{L}(f(\mathbf{h}, \mathbf{r}, \mathbf{t}), f(\tilde{\mathbf{h}}, \mathbf{r}, \tilde{\mathbf{t}})), \quad (3)$$

where t is a training step, ∇_{θ} is the gradient with respect to θ and α is the learning rate. We outline the algorithm for training BioBLP in Algorithm 1.

Input: Knowledge graph (V, R, E, D) with entity-to-attribute map $d : V \rightarrow D$
 Modality mapping $\text{modality_of} : D \rightarrow \{0, 1, \dots, k\}$
 Embedding functions of entities and relations: $\{e_0, \dots, e_k\}, e_r$
 Scoring function f
 Loss function \mathcal{L}
 Number of steps T , Learning rate α

Function $\text{emb}(\text{entity } x)$:

```

 $m \leftarrow \text{modality\_of}(d(x));$ 
if  $m = 0$  then
   $x = e_0(x)$ ;
else
   $x = e_m(d(x))$ ;
end
return  $x$ ;
```

$\theta^{(0)} \leftarrow$ parameters in $\{e_0, \dots, e_k\}, e_r$;

```

for  $t = 1 \dots T$  do
  for  $(h, r, t) \in E$  do
     $(\tilde{h}, r, \tilde{t}) \leftarrow \text{corrupt}(h, r, t);$ 
     $\theta^{(t)} \leftarrow \theta^{(t-1)} - \alpha \nabla_{\theta} \mathcal{L}(f(\text{emb}(h), e_r(r), \text{emb}(t)), f(\text{emb}(\tilde{h}), e_r(r), \text{emb}(\tilde{t})))$ 
  end
end
return  $\theta^{(T)}$ 
```

Algorithm 1 BioBLP: Learning embeddings in a KG with entity attributes

Implementing BioBLP

The modular formulation of BioBLP is general enough to enable experiments with different choices for each of the modules, and presents an opportunity for incorporating domain knowledge into the design of different attribute

encoders. In this section we describe the particular choices that we make for training BioBLP with the BioKG graph.

Considering that we retrieve attribute data for proteins, drugs, and diseases in BioKG, we define encoders for these three types. Recent works have demonstrated the effect of using pretrained embeddings in link prediction models that use textual descriptions [23, 32, 33]. Pretraining is done on a task for which a large dataset is available, such as masked language modeling [35]. Intuitively, the use of large pretrained models provides a better starting point for training, compared to randomly initializing an encoder and training for link prediction. In the context of the biomedical domain, we explore this strategy by considering three pretrained models as attribute encoders: ProtTrans [55], MolTrans [56], and BioBERT [57].

Protein encoder ProtTrans is a model for encoding amino acid sequences of proteins, where each of the input tokens is an amino acid [55]. The output is a matrix of “contextualized” amino acid embeddings of size $n_e \times l$, where l is the length of the sequence and n_e the dimension of the ProtTrans embeddings. Encoding proteins with ProtTrans in BioKG is challenging: we found that some sequences can have lengths up to 35,213, which greatly increases memory requirements. To reduce memory usage during training, we choose not to optimize the parameters of the ProtTrans encoder. Given the output matrix of amino acid embeddings, we obtain a single protein embedding by taking the average over the sequence, and then linearly projecting the result to the embedding space of dimension n via a matrix of size $n \times n_l$ with learned parameters. The reason for averaging the embeddings is that due to computation constraints, it is not feasible to fine-tune the ProtTrans encoder. If fine-tuning the encoder were feasible, it would be possible to use the embedding of a special marker at the beginning of the sequence. We thus use the average as an aggregated representation.

Molecule encoder MolTrans is a model initially trained to learn representations of molecules, by translating them from their SMILES representation to their IUPAC name [56]. We use it to obtain embeddings for molecules in BioKG for which we have a SMILES representation. The output of MolTrans is a matrix of embeddings of size $n_e \times l$, where l is the length of the SMILES string and n_e is the size of the MolTrans embeddings. Similar to the case of the protein encoder, we do not optimize the parameters of MolTrans due to memory constraints. Instead, we apply a layer of self-attention [58] and select the output embedding at the first position, which corresponds to a special marker for the beginning of the sequence. This embedding is then linearly projected to an embedding of size n , and the result is used as

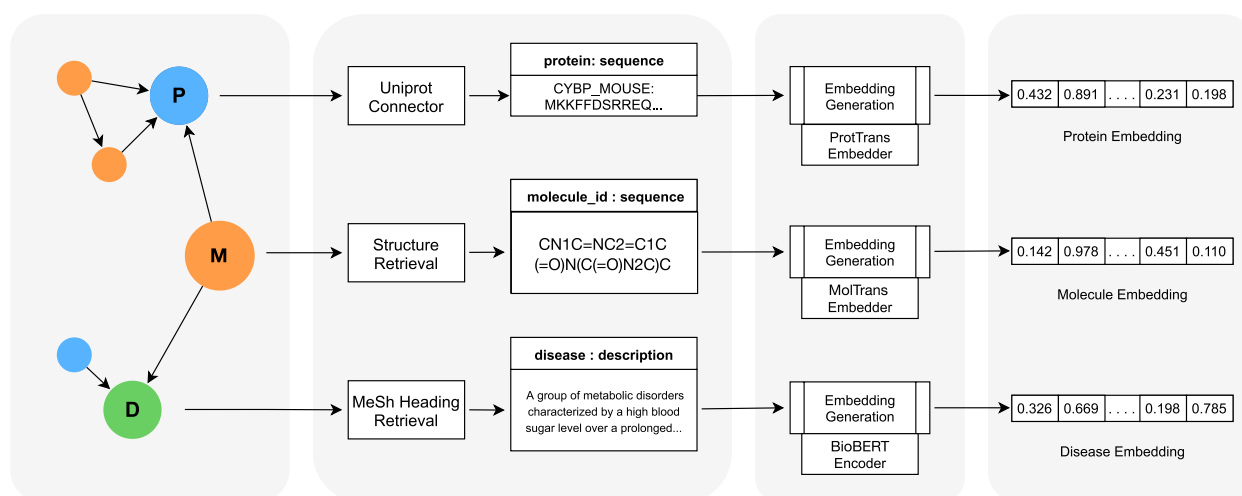


Fig. 3 The architecture of BioBLP for the BioKG graph. On the left we illustrate the graph with the three entity types that have attribute data: proteins (P), molecules (M, representing drugs), and diseases (D). We omit entity types without attributes for clarity. For the three types, we retrieve relevant attribute data, which are the inputs to encoders for each modality. The output of the encoders is a vector representation of the entities

the embedding of the molecule. In the case of the molecule encoder, the learnable parameters are the weights in the self-attention layer, and the linear projection matrix.

Disease encoder We rely on BioBERT [57] to encode textual descriptions of entities. BioBERT is a language model trained to predict words that have been removed from some input text, which is also known as the masked language modeling task [35]. In contrast to models like BERT [35] that have been trained over diverse sources from the internet, BioBERT has been pretrained with scientific articles obtained from PubMed and PMC. Given an input textual description, we select the output embedding of the special marker at the start of the description, and then project it to an embedding of size n . As opposed to the protein and molecule encoders, we optimize all the layers in the BioBERT encoder as well as the projection layer. We can afford this by limiting the length of the descriptions to a maximum of 64 tokens, which helps to curb the memory requirements during training. Previous work has shown that further increasing the length of text encoders for embedding entities in KGs brings diminishing returns [23].

Choice of pretrained encoder models. An important tradeoff between computational complexity and accuracy arises when selecting encoders. While state-of-the-art methods can provide better representations [59], we prioritized publicly available encoders that would fit within our computation budget.

In summary, we provide an implementation of BioBLP for the BioKG graph that uses encoders of amino acid sequences for proteins, SMILES representations for molecules, and textual descriptions of diseases. Following the

BioBLP framework, other entity types like pathways and gene disorders are embedded using lookup table embeddings. We illustrate the architecture in Fig. 3.

Previous research has shown that in spite of continuous developments in the area of KG embeddings, the choice of scoring function is a matter of experimentation and the results are highly data-dependent [23, 50, 51]. For these reasons, we experiment with three scoring functions: TransE [18], ComplEx [19], and RotatE [20]. For the loss functions, we consider the margin ranking loss, the binary cross-entropy loss, and the cross-entropy loss [51]. Our implementation uses the PyKEEN library [60] and is publicly available⁶.

Efficient model pretraining

Training KG embedding methods that use encoders of attribute data has been shown to significantly increase the required computational resources for training [23, 32, 33]. Furthermore, BioBLP consists of a set of attribute encoders, together with lookup table embeddings for entities with missing attributes. Training such a model from scratch requires all the encoders to adapt together for the link prediction objective. If lookup table embeddings are initialized randomly as is customary, attribute encoders will adjust to their random-level performance, greatly slowing down the rate of learning. We thus propose the following strategy to overcome this problem:

⁶ Available at <https://github.com/elsevier-AI-Lab/BioBLP>.

1. Train a BioBLP model with lookup table embeddings for *all* the entities in the graph.
2. Train a BioBLP model with encoders for entities with attribute data, and for entities without such data, initialize their lookup table embeddings from the model trained in step 1.

Step 1 is much cheaper to run since it does not make use of any attribute encoders. This speeds up training, and facilitates hyperparameter optimization. In step 2, the encoders no longer have to adapt to randomly initialized lookup table embeddings, but to pretrained embeddings that already contain useful information about the structure of the graph, which reduces the number of training steps required.

Answering RQ1 Given the question, “How can we incorporate multimodal attribute data into KGE models under a common learning framework?”, we find that

- A survey of the literature shows that existing methods ignore attribute data, or assume a uniform modality (such as text) for all entities.
- We design BioBLP as a method that enables the study of the effect of incorporating multimodal data into knowledge graph embedding, under a common learning framework where all encoders are optimized jointly.
- We propose a strategy for efficient model pretraining that improves the practical application of BioBLP.

Evaluation

In our experiments, we aim to answer research questions RQ2 and RQ3, related to how the different attribute encoders in BioBLP compare with traditional KG embedding methods that rely on the structure of the graph only.

We focus on two tasks to answer this question: link prediction, and the DPI benchmark present in BioKG.

Link prediction

The link prediction evaluation is standard in the literature of KG embeddings [50, 51]. Once the model is trained, the evaluation is carried out over a held-out test set of triples that is disjoint from the set of triples used during training.

For a triple (h, r, t) in the test set, a list of scores for *tail prediction* is obtained by fixing h and r and computing the scoring function $f(\mathbf{h}, \mathbf{r}, \mathbf{x})$, for all possible entity embeddings \mathbf{x} . The list of scores is sorted from high to low, and the rank of the score for the true tail entity t is determined. We repeat this for all triples in the test set, collecting a list of ranks for tail prediction. For head

prediction, the process is similar, but instead we fix r and t and collect the ranks of the true head entity h given scores for the head position. We collect the ranks for head and tail prediction in a list (p_1, \dots, p_m) . From this list we compute the Mean Reciprocal Rank, defined as

$$\text{MRR} = \frac{1}{m} \sum_{i=1}^m \frac{1}{p_i}, \quad (4)$$

and Hits at k , defined as

$$\text{H@k} = \frac{1}{m} \sum_{i=1}^m \mathbb{1}[p_i \leq k], \quad (5)$$

where $\mathbb{1}[\cdot]$ is an indicator function equal to 1 if the argument is true and zero otherwise.

It is important to note that by averaging over the set of evaluation triples, these metrics are more influenced by high-degree nodes, simply because they have more incident edges. This means that it is possible that a model can score highly in the link prediction task by correctly predicting edges for a small number of entities that account for a significant number of the triples in the graph [61].

We use the MRR and H@k metrics to report the performance of the embeddings obtained with BioBLP. As baselines that do not incorporate attribute data, we consider TransE, ComplEx, and RotatE. For a direct comparison, we experiment with variants of BioBLP that use the same scoring functions as in these baselines. This allows us to determine when improvements are due to the choice of scoring function, rather than a result of incorporating attribute data.

Hyperparameter study Previous research has shown that KG embedding models and the optimization algorithms used to train them can be sensitive to particular choices of hyperparameters [50], including in the biomedical domain [52]. Hence, a key step in determining the performance of our baselines against models that incorporate attributes is an exhaustive optimization of their hyperparameters. For each of the three baselines we consider, we carry out a hyperparameter study with the aim of optimizing their performance before comparing them with BioBLP. We focus on four hyperparameters as listed in Table 3, where we also describe the intervals and values we considered. For each baseline, we used Bayesian optimization to search for hyperparameters, running a maximum of 5 trials in parallel. We ran a total of 90 different hyperparameter configurations with TransE, 180 with ComplEx, and 180 with RotatE. For TransE we only consider the margin ranking loss function, while for ComplEx and TransE we experiment with the binary cross-entropy and cross-entropy, which leads to twice the number of experiments.

Table 3 Intervals and values used in the hyperparameter study we carried out to optimize the TransE, ComplEx, and RotatE baselines

Hyperparameter	Values
Learning rate	Log-uniformly sampled from $[1 \times 10^{-3}, 1.0]$
Regularization weight	Log-uniformly sampled from $[1 \times 10^{-6}, 1 \times 10^{-3}]$
Batch size	{128, 256, 512, 1024}
Loss function	Margin ranking, binary cross-entropy, cross-entropy

Benchmarks

We evaluate the utility of entity embeddings by testing their performance as input features for classifiers trained on related biomedical tasks. We focus on the DPI benchmark obtained from BioKG. As it is common in literature, we interpret this task as binary classification [62–64]. We use the pairs in the benchmarks as positive instances to train the classifiers. A common issue in modelling efforts in biomedical domain is the lack of true negative samples in public benchmark datasets. Therefore, we sample negatives from the combinatorial space assuming a pair that is unobserved in the dataset is a candidate negative sample [39, 62–64]. Furthermore, we filter out any triples appearing in BioKG from the list of candidate negatives. We use a ratio of positive to negatives of 1:10.

For a given pair of entities, we concatenate their embeddings into a single vector that is used as input to a classifier. We evaluate a number of baselines for computing embeddings, including random embeddings, entity embeddings that do not rely on the graph, and KG embeddings. We compare these to embeddings obtained with BioBLP. In particular, we use the following embeddings:

- **Random.** Every entity is assigned a randomly generated feature vector from a normal distribution. We use this method to determine a lower bound of chance performance. Each entity is assigned an 128-dimensional embedding.
- **Structural.** Drug entities are represented by their MolTrans embeddings of 512 dimensions, and proteins are represented by their ProtTrans embedding, averaged across tokens resulting in a 1024-dimensional vector.
- **KG.** Entities are represented by TransE, ComplEx, and RotatE embeddings trained on BioKG. The entity embeddings from these models all have 512 dimensions.
- **BioBLP.** Entities are represented from the attribute encoders in BioBLP models trained using the RotatE

scoring function, on the BioKG graph. The entity embeddings from the BioBLP-D, BioBLP-M and BioBLP-P models have 512 dimensions.

When using Structural Embedding, it may occur that a particular entity does not have a corresponding feature, for example, because we are missing the chemical structure embeddings for a set of molecules (as detailed in Table 1) because these are biologic therapies and do not have a SMILES property. In these cases, we impute the missing embedding by the mean of all entity embeddings for that entity type.

Training and evaluation We experiment with Logistic Regression (LR), Random Forest (RF) and Multi Layer Perceptron (MLP) models for classification. To aid the training of the classifiers in the imbalanced data setting, we use class weights inversely proportional to the class frequency in the LR and RF models, and undersample the majority class in the training of the MLP models.

The performance estimates of the classification models is based on 5-fold stratified cross-validation. The folds are kept consistent for training and testing across all models. For each fold, hyperparameter optimization is performed to establish optimal model parameters on the training split. Ten percent of the training split is reserved for validation in the tuning procedure. We use Optuna to perform the hyperparameter search, which uses a tree-structured Parzen estimator algorithm for sampling the space of hyperparameters. For all models, 10 trials are performed to find the best parameters on the validation set that maximize performance on the area under precision-recall curve (AUPRC) and area under ROC curve (AUROC) metrics. From these trials the best parameters are determined and the model is retrained on the full training set and scored against the test set.

Metrics In the evaluation procedure of the classification models we use Precision, Recall, F1, AUROC and AUPRC values. The AUROC is commonly reported in bioinformatics literature but is argued to be an overly optimistic measure of performance [65]. For this reason, we additionally report the AUPRC because this metric is particularly informative in imbalanced datasets due to its sensitivity to true positives among all positive predictions.

Results

In this section we describe our findings on the link prediction performance of embedding methods that do not use attributes, and the effect of incorporating them with BioBLP. We additionally present the results on the DPI benchmark, and we then provide a fine-grained analysis of the performance provided by BioBLP.

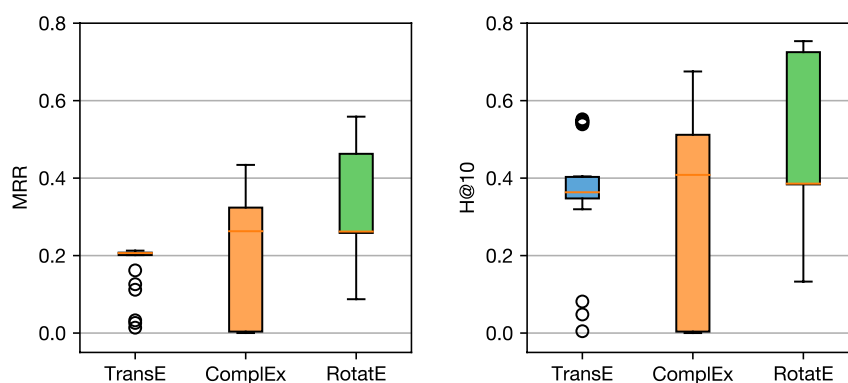


Fig. 4 Distribution of MRR (left) and H@10 (right) performance metrics for link prediction, across all values explored during hyperparameter optimization of the baseline models. TransE performs rather consistently, but poorly overall. ComplEx results on a broad range of outcomes, indicating that it is very sensitive to hyper-parameters. RotatE performs similar to ComplEx when comparing their medians, but has a smaller spread and with specific parameters it results in the best performance

Hyperparameter study

We illustrate in Fig. 4 the distribution of link prediction metrics for the three baselines that we consider, across all the configurations explored during the hyperparameter study, evaluated on the test set of triples. These results confirm the importance of thoroughly exploring the space of hyperparameters, as the three baselines exhibit high variance in the final performance. For example, we observe that for specific configurations, TransE can outperform RotatE, even though by design RotatE is more expressive at modeling different kinds of relations [20].

Across all the hyperparameter trials that we run, we observe that RotatE yields consistently better results, with the worse RotatE configuration outperforming multiple configurations of TransE and RotatE. We select the best configurations of these three models based on the validation set performance, and we use them in the next set of experiments.

Link prediction performance

We present link prediction results of BioBLP under different configurations, corresponding to variants along three axes: the scoring function, the type of encoder, and whether or not the model includes our proposed pre-training strategy. For scoring functions, we experiment with TransE, ComplEx, and RotatE, and for encoders, we experiment with Proteins (P), Molecules (M), and Diseases (D). An example of a specific configuration for a BioBLP model that uses the TransE scoring function and implements a molecule encoder is TransE+BioBLP-M.

The results are shown in Table 4, where we compare BioBLP with our baselines. We observe that in all cases, introducing different kinds of attribute encoders results in lower performance when compared to the baselines that rely on the structure of the graph alone to learn

embeddings. However, we note that when combined with the RotatE scoring function, BioBLP yields competitive performance when using protein and disease encoders. This implies a trade-off to be considered when learning embeddings from attribute data: while lookup table embedding methods are effective at link prediction, they are limited to the set of entities observed during training. BioBLP results in a drop in performance, but it provides a more general embedding method that can generate representations for entities outside the training set of entities.

As observed during the hyperparameter study, RotatE results in the best performance across all models. Its performance also transfers to BioBLP models that use the RotatE scoring function, indicating that regardless of the mechanism used for embedding entities (be it lookup tables, or attribute encoders), the geometry of the space induced by RotatE and its scoring function are particularly effective for link prediction on BioKG. When using BioBLP with other scoring functions, performance is significantly lower. Notably, ComplEx+BioBLP-D fails at solving the link prediction task.

When focusing on the best-performing group of models that use RotatE, we can see that the lowest performance is obtained when using molecule encoders, suggesting that further improvements are required in how molecules are represented and encoded in BioBLP-M models. When using protein and disease encoders, link prediction performance increases, with RotatE+BioBLP-D yielding the highest performance over all our models that use encoders.

We note that our proposed strategy for pretraining always increases the performance over the corresponding models that have not been pretrained. The magnitude of such improvements varies, with some resulting in slightly

Table 4 Link prediction performance on the BioKG dataset (in percent)

Method	Pretrained	MRR	H@1	H@3	H@10
TransE		20.61	12.15	22.14	36.27
+ BioBLP-P		14.76	8.42	15.60	25.92
+ BioBLP-P	✓	14.87	8.47	15.67	26.16
+ BioBLP-M		6.56	4.46	6.78	9.90
+ BioBLP-M	✓	8.72	6.22	9.42	13.10
+ BioBLP-D		7.42	4.91	7.39	11.62
+ BioBLP-D	✓	15.74	11.73	16.44	23.04
ComplEx		42.75	31.55	48.09	65.50
+ BioBLP-P		16.53	11.13	17.17	26.92
+ BioBLP-P	✓	34.88	25.11	39.52	54.83
+ BioBLP-M		1.68	1.24	1.86	2.49
+ BioBLP-M	✓	1.78	1.37	1.93	2.54
+ BioBLP-D		0.01	0.00	0.00	0.02
+ BioBLP-D	✓	0.45	0.36	0.43	0.55
RotatE		55.20	44.46	61.95	74.76
+ BioBLP-P		45.29	35.60	51.33	62.89
+ BioBLP-P	✓	47.30	36.56	54.59	66.10
+ BioBLP-M		10.40	7.02	10.40	16.30
+ BioBLP-M	✓	14.34	11.14	14.79	19.78
+ BioBLP-D		11.60	8.79	12.43	16.67
+ BioBLP-D	✓	49.68	40.62	55.30	66.26

higher performance and others displaying pronounced differences. The best variant that we obtained corresponds to a pretrained RotatE+BioBLP-D model, whose performance increases by approximately four times when compared to the model that is not pretrained. We test

for significance using a paired Welch’s *t*-test by comparing the reciprocal rank per triple in the test set, between models trained from scratch and pretrained models. We illustrate the impact of pretraining in Fig. 5, where we plot the increase in MRR, and the performance until convergence as a function of time. With the exception of BioBLP-P with the TransE scoring function, we find that the increase in MRR provided by pretraining is statistically significant ($p < 0.05$). We also note that pretraining reduces the time required to converge and reach higher link prediction performance, compared to no pretraining. Even when allocating more time resources to the setting without pretraining, the model reaches much lower values of MRR.

Benchmarks

We present results on the DPI benchmarks using the BioBLP models that employ the RotatE scoring function, since these provided significantly higher performance in the link prediction task compared to variants using other scoring functions.

In Figure 6 we report the AUROC and AUPRC metrics for the RF classifiers, which resulted on better performance compared to LR and MLP classifiers, on the five test folds for positive-negative ratio 1:10. We focus on these models to discuss the impact of the type of embedding for solving the DPI task. We provide extended results for all classifiers in Appendix A.

We observe that classifiers trained on BioBLP embeddings remain competitive with the baseline embeddings considered in this work. The Random baseline (not

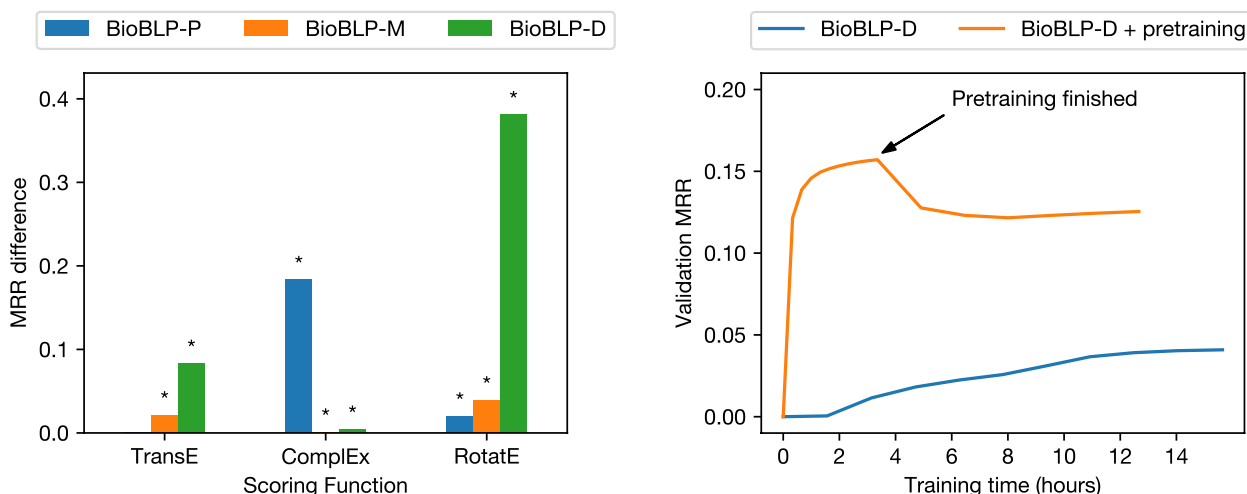


Fig. 5 The impact of our proposed pretraining strategy. Left: increase in MRR due to the pretraining strategy over base models trained from scratch. Asterisks indicate significance at p -values less than 0.05. Right: A comparison of the runtime required until convergence on the validation set, when training BioBLP-D from scratch, and when including pretraining. During pretraining, attributes are ignored, and after pretraining we start optimizing the attribute encoders

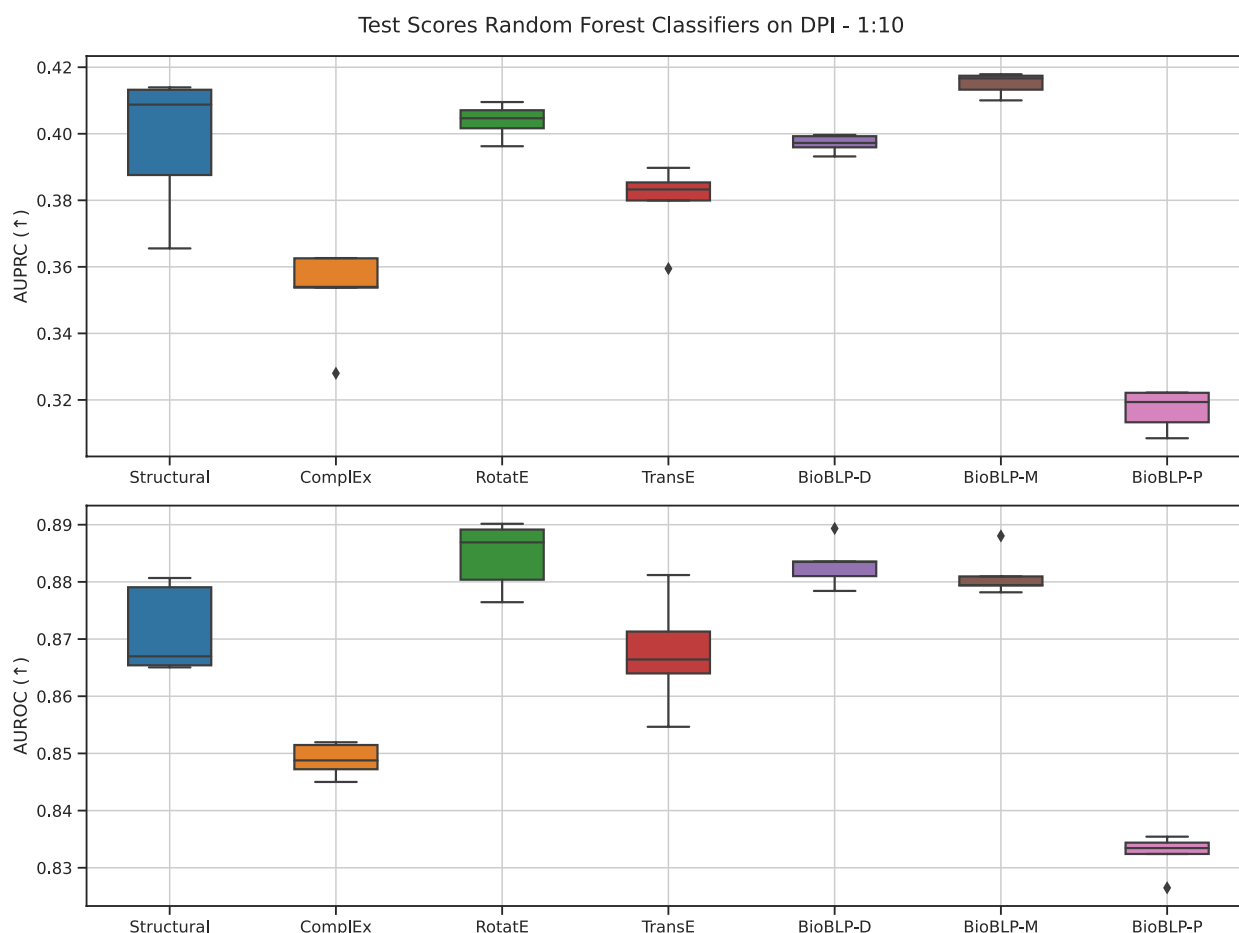


Fig. 6 Classification performance for Random Forest Classifiers for DPI benchmark with positive-negative ratio of 1:10. The box for the baseline classifier using Random embeddings is omitted for clarity, reaching mean AUPRC of 0.192 and mean AUROC of 0.720. The median is indicated by the horizontal line, the quartiles indicated by the box, the min and max by the whiskers and individual points are outliers

shown in Fig. 6) reached mean AUROC of 0.720 and mean AUPRC of 0.192. All models score above these values, confirming the classifiers learn the task better than chance when given the proposed embeddings as features. Among the baselines, RotatE and Structural embeddings lead to highest mean scores in the AUROC and AUPRC metrics. The model trained on BioBLP-M reaches highest mean AUPRC overall, but on the AUROC metric scores below BioBLP-D and RotatE embeddings. Notably, the models trained on BioBLP-P embeddings score lowest in this setup, while this model was among the top scoring models in the link prediction setup.

These results show that embeddings from BioBLP-M and BioBLP-D, which encode molecular structures and disease descriptions, provide competitive performance when compared to lookup table embeddings. In contrast, embeddings from BioBLP-P, which encodes protein sequences, perform worse among other types of embeddings. In particular, BioBLP-P performs worse than the

Structural baseline, whose embeddings are obtained from ProtTrans and MolTrans. Since we also employ ProtTrans embeddings as inputs for BioBLP-P, this indicates that using BioBLP-P *degrades* the quality of the ProtTrans embeddings, thus further work is required in properly incorporating protein sequences into encoders for BioBLP. We do not observe such degradation in BioBLP-M, which uses MolTrans embeddings as inputs but outperforms the Structural baseline.

It is notable that we find BioBLP-D among the top performing models in this task involving molecule and protein entities. It suggests that the embeddings for these entity types receive meaningful signals from the disease attribute encoding during training on the link prediction objective, through their relations with disease entities.

Answering RQ2 We next summarize our answers to the question: “What is the performance in comparison with baselines that do not incorporate attribute data

on link prediction and drug-protein interaction benchmarks?” Our experiments show that:

- When their hyperparameters are properly tuned, baselines not incorporating attribute data can exhibit higher link prediction performance.
- The best performing variants of BioBLP make use of protein and disease encoders, and the RotatE scoring function. It results in link prediction performance that though lower than RotatE, remains competitive given the complexity of the task. The fact that molecule encoders underperform indicate that further work is required in the respective encoder to properly capture the relevant features of molecules, and alternatives to using MolTrans embeddings should be considered.
- These results extend to the DPI benchmark. Furthermore, we find that the disease encoder has an indirect positive effect in the benchmark (which does not contain diseases), since diseases are linked to proteins and molecules in the graph.

Performance analysis

Lookup table embeddings such as RotatE have been observed to suffer from node-degree bias, obtaining high link prediction scores by learning to predict links for a small fraction of entities that are responsible for a large number of triples in the graph [61]. These models show poor link prediction performance in sparser regions of the graph comprised by entities with a low degree of incoming or outgoing edges.

We investigate the performance of link prediction models with respect to the degree of the entities being predicted, focusing on a comparison between RotatE and BioBLP-D, which uses an encoder of descriptions of diseases. We select all test triples involving diseases, and we compute the difference in MRR between BioBLP-D and RotatE when i) given a non-disease entity and a relation, *diseases are predicted*; and ii) given a disease entity and a relation, *non-disease entities are predicted*. A positive difference indicates that RotatE results in higher performance.

The results are shown in Fig. 7, where we present the difference in MRR, together with a visualization of the degree distribution for predicted entities. From Fig. 7, we observe that BioBLP-D outperforms RotatE when the predicted entity is a disease, over a subset of low degree entities that range from a degree of 1 to approximately 20. For entities with a degree between 20 and 100 we observed mixed results, with multiple cases in which BioBLP-D produces better results. In the case of

high degree entities, RotatE consistently performs better than BioBLP-D when predicting a disease.

When given a disease and predicting a non-disease, Fig. 7 shows that RotatE works better across different degree values, with the exception of a small set of high-degree diseases where BioBLP-D results in higher MRR.

The degree distributions visualized in Fig. 7 suggest that a substantial amount of entities in the graph have low to medium degrees, while entities of high degree are few. This is expected in graphs where the distribution of the degree follows a power law, and it can occur in the biomedical domain where some entities are involved in multiple processes, while other entities have specific functions or are understudied [66]. We confirm this by examining the quantiles of degree values for different entity types in Table 5. We note that 50% of the entities in the graph have degrees of up to 6, which is well below the degree values at which BioBLP-D outperforms RotatE. Interestingly, this result is not apparent in the standard link prediction evaluation, where metrics are micro-averaged over the complete set of triples and are thus dominated by the performance over high-degree entities.

As observed in previous work [61], relying heavily on such micro averages might bias the selection of models that perform well only on frequent entities, which occur as high degree nodes in the graph. This motivates us to analyze a *macro* version of MRR, in which we compute the average MRR for each value of node degree. The result is an average MRR where the performance over each node degree value contributes equally, rather than being dominated by high degree nodes as in the case of the standard, micro average MRR.

In Fig. 8 we present the difference in macro MRR between RotatE and BioBLP as we consider cumulative values of degree: for each value d of degree in $1, 2, \dots$, we compute the difference in macro MRR for all nodes with degree less than or equal to d . The result shows that there is a cutoff point below which BioBLP-D outperforms RotatE at a degree value of 74. Table 5 shows that this value accounts for more than 75% of the diseases in the graph. This implies that applications concerned with low degree entities would benefit from BioBLP-D once this cutoff has been determined. For BioKG, this cutoff comprises the majority of the diseases in the graph.

Answering RQ3 Based on our experimental results, we now answer the question: “Under what conditions do we observe positive or negative effects of incorporating attribute data using our proposed framework?”. We find that:

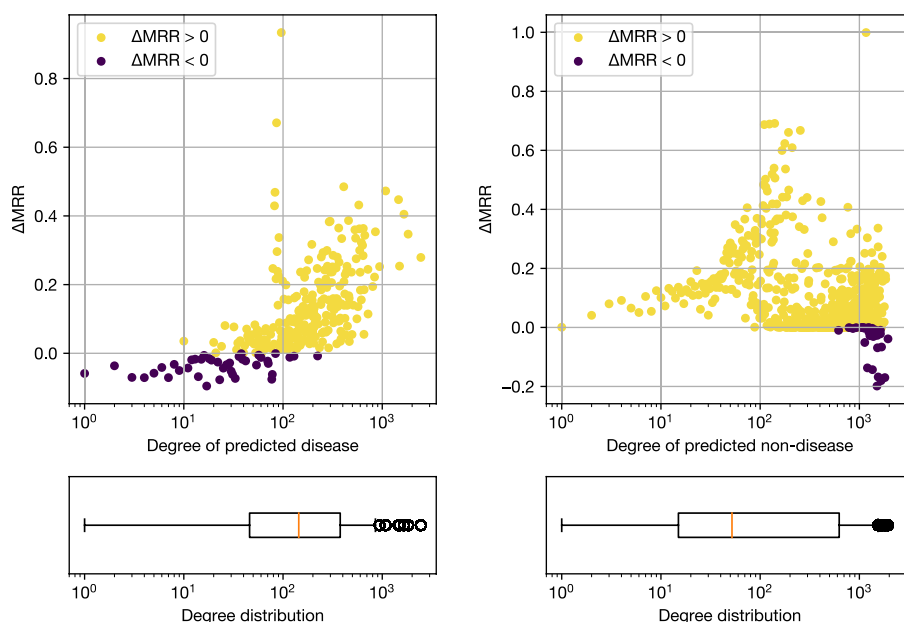


Fig. 7 Difference in Average MRR between RotatE and BioBLP-D predictions against node degree of *disease* (left) or *non-disease* (right) being predicted, together with the degree distribution of predicted entities. A negative ΔMRR means that BioBLP-D outperforms RotatE, while a positive value means that RotatE has higher performance

Table 5 Quantiles for node degree per entity type in the BioKG training graph used in our work. The node degree of an entity is the sum of the incoming and outgoing edges from the entity

Statistic	mean \pm std	min	25%	50%	75%	max
Proteins	9.60 \pm 19.33	1	2	4	9	480
Diseases	25.89 \pm 85.37	1	3	5	15	2448
Drugs	218.85 \pm 351.24	1	2	6	369	1942
All	27.95 \pm 119.00	1	2	4	10	2448

- When incorporating attribute data, the choice of the scoring function is crucial for performance. Using BioBLP with the RotatE scoring function results in significantly better performance in comparison with using TransE or ComplEx as the scoring function.
- The observed high performance of RotatE can be partially attributed to the fact that test triples are dominated by high-degree entities, where BioBLP underperforms.
- When the focus is in low-degree entities BioBLP can have a positive effect in comparison with RotatE. When values of node degree contribute equally to the computation of MRR, we find a threshold of node degree below which BioBLP outperforms RotatE. This threshold covers more than 75% of the diseases in the graph.

Discussion

With BioBLP, we present a framework for studying the effect of incorporating multimodal attribute data into KG embeddings. Doing so requires reconsidering the requirements of existing work on embedding biomedical KGs that assumes that all entities in the KG are associated with the same data modality or that such information is available for all entities in the graph. We observe that when combined with the RotatE scoring function, BioBLP-D provided competitive, yet lower performance in the link prediction and DPI benchmark, in comparison with a RotatE baseline that uses lookup table embeddings and thus does not take attribute data into account. This is in line with prior works that explore the use of textual descriptions to learn entity embeddings [23]. Optimizing the molecule and protein encoders in BioBLP-M and BioBLP-P proved more challenging, resulting in

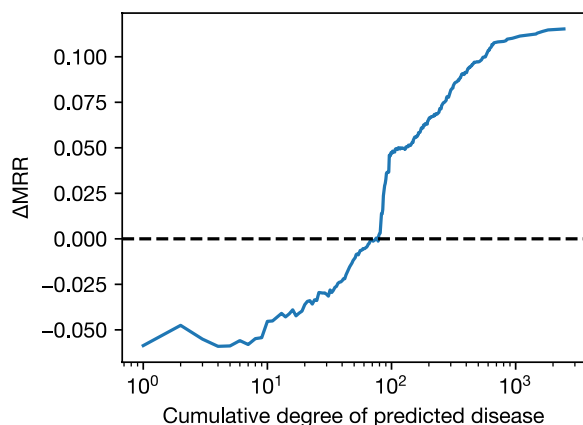


Fig. 8 Difference in *macro* MRR between RotatE and BioBLP-D when a disease is predicted, as we consider increasing values of node degree. We observe that when considering nodes with a degree of 74 or less, BioBLP-D results in improved link prediction performance (shown as negative values of Δ MRR)

embeddings with lower performance. Future work could attempt to find ways to create encoders for these modalities that are expressive while keeping training over large KGs scalable.

On one hand, these findings indicate that RotatE is an effective model for link prediction on BioKG, as well as a suitable scoring function when used in the BioBLP framework. However, we note that the common evaluation protocol where link prediction metrics are micro-averaged over a test set of triples biases the results such that the average is dominated by high-degree entities in the graph. This issue has been highlighted in prior work on evaluation protocols for link prediction methods [61]. We argue that for tasks of scientific discovery, the performance over low degree entities is essential - if not more important - since these are understudied areas of the graph that would benefit more from predicting potential links that do not exist in the KG yet. Our analysis shows that BioBLP-D achieves better performance than RotatE when predicting edges to low degree diseases, which comprise a substantial amount of the complete set of entities in the KG. Based on this finding, we call for future scientific discovery work which develops alternative evaluation protocols for assessing link prediction methods in that context.

Further, we observed that training BioBLP results in increased computational costs, since the attribute encoders are deep neural networks with multiple layers and more parameters in comparison with the shallow, lookup table embeddings found in models like RotatE. Our results show that pretraining is an effective strategy for lowering the computational cost of training BioBLP.

Pretraining is done using a model that is drastically cheaper to optimize, such as RotatE, so that once the relevant modules are transferred to BioBLP, less training steps are required and the resulting performance is higher compared to omitting pretraining.

A useful property of BioBLP is its ability to encode arbitrary entities, rather than entities predefined in the training set, as long as attribute data is available for them. While our focus was comparing lookup table embeddings with BioBLP, future work could focus on the inductive setting where lookup table embeddings are not applicable. Our framework could also be further extended by considering that a single entity might have not one but multiple attributes, such as SMILES representations *and* molecular graphs.

The use of attribute encoders in BioBLP suggests a new interface for leveraging the information contained in a KG with attribute data. One interesting application for our proposed approach is the development of a natural language interface with the KG for biomedical tasks. Using the text encoder component, we have the ability to make predictions such as the effectiveness of a specific drug for a particular disease, even if the disease is not explicitly present in the KG.

Another use case that we would explore in future work is expanding the KG with new entities. As KGs are inherently incomplete and require frequent updates, we could use BioBLP to compute embeddings for incoming entities, discover the closest entities by attribute similarity, and predict the most appropriate links with entities already in the KG.

Conclusions

In this work, we have studied the effect of incorporating multimodal attribute data into biomedical KG embeddings. To this end, we have presented BioBLP, a modular framework for learning embeddings of entities in a KG with attribute data, that also supports entities with missing attribute data. We have additionally introduced a pre-training strategy for reducing training runtime.

Experiments on general link prediction show that BioBLP does not yet match baselines that do not use attribute data, though it performs closely. In drug-protein interaction benchmarks, its performance matches the baselines. An in-depth analysis reveals that BioBLP in fact outperforms the baselines over entities below a certain node degree. This subset of entities covers at least 75% of the diseases in the graph. We argue that these low node-degree entities are of special interest in the area of scientific discovery.

In future work, we foresee the investigation of the tradeoff that arises when designing attribute encoders that are expressive yet scalable, as well as applications

enabled by the use of multimodal attribute encoders in biomedical knowledge graphs, including predictive problems over entities not seen during training.

Appendix A additional benchmark results

We present extended results on the DPI benchmark for all three classifiers (logistic regression, multi-layer perceptron, and random forest) on Table 6. In most cases, we observe that the random forest outperforms the other classifiers on the DPI task when using KG embeddings as inputs.

Table 6 Classification performance on DPI benchmark with ratio 1:10. Reported are the mean (std) over the 5 best models scored on the test folds. Values in bold indicate the highest metric for each feature type

Feature	Classifier	DPI-FDA		
		AUPRC	AUROC	F1
Random	LR	0.129 ± 0.003	0.597 ± 0.001	0.602 ± 0.010
	MLP	0.139 ± 0.006	0.635 ± 0.014	0.594 ± 0.004
	RF	0.192 ± 0.004	0.720 ± 0.005	0.889 ± 0.001
Structural	LR	0.153 ± 0.003	0.690 ± 0.010	0.617 ± 0.010
	MLP	0.244 ± 0.006	0.771 ± 0.003	0.665 ± 0.010
	RF	0.398 ± 0.021	0.871 ± 0.008	0.885 ± 0.016
ComplEx	LR	0.170 ± 0.002	0.706 ± 0.004	0.639 ± 0.004
	MLP	0.339 ± 0.006	0.838 ± 0.004	0.755 ± 0.003
	RF	0.352 ± 0.014	0.849 ± 0.003	0.892 ± 0.008
RotatE	LR	0.180 ± 0.002	0.722 ± 0.005	0.643 ± 0.003
	MLP	0.453 ± 0.008	0.886 ± 0.002	0.794 ± 0.006
	RF	0.404 ± 0.005	0.885 ± 0.006	0.905 ± 0.006
TransE	LR	0.170 ± 0.003	0.691 ± 0.021	0.650 ± 0.017
	MLP	0.350 ± 0.009	0.846 ± 0.005	0.750 ± 0.003
	RF	0.380 ± 0.012	0.868 ± 0.010	0.894 ± 0.015
BioBLP-D	LR	0.177 ± 0.003	0.723 ± 0.003	0.641 ± 0.002
	MLP	0.447 ± 0.008	0.885 ± 0.003	0.792 ± 0.002
	RF	0.397 ± 0.003	0.883 ± 0.004	0.907 ± 0.004
BioBLP-M	LR	0.166 ± 0.002	0.713 ± 0.002	0.634 ± 0.003
	MLP	0.401 ± 0.007	0.865 ± 0.002	0.764 ± 0.002
	RF	0.415 ± 0.003	0.881 ± 0.004	0.895 ± 0.006
BioBLP-P	LR	0.170 ± 0.006	0.683 ± 0.005	0.636 ± 0.005
	MLP	0.343 ± 0.008	0.820 ± 0.005	0.744 ± 0.015
	RF	0.317 ± 0.006	0.832 ± 0.004	0.899 ± 0.002

Abbreviations

AUPRC	Area Under the Precision-Recall Curve
AUROC	Area Under the Receiver Operating Characteristic
DTI	Drug-Target Interaction
H@k	Hits at k

KG	Knowledge Graph
LR	Logistic Regression
LS	Life Sciences
MLP	Multi-Layer Perceptron
MRR	Mean Reciprocal Rank
RF	Random Forest

Acknowledgements

Not applicable.

Authors' contributions

D.D., D.A., P.M., and T.P. prepared and analyzed the data, designed models and experiments, and executed experiments. M.C. and P.G. supervised the project. All authors read and approved the manuscript.

Funding

This project was funded by Elsevier's Discovery Lab.

Availability of data and materials

The datasets generated and/or analysed during the current study are available in the Zenodo repository, available at <https://doi.org/10.5281/zenodo.8005711>. Our software implementation is open source and available at <https://github.com/elsevier-AI-Lab/BioBLP>.

Declarations

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Competing interests

The authors declare no competing interests.

Author details

¹Vrije Universiteit Amsterdam, Amsterdam, The Netherlands. ²University of Amsterdam, Amsterdam, The Netherlands. ³Elsevier B.V., Amsterdam, The Netherlands. ⁴Discovery Lab, Elsevier, Amsterdam, The Netherlands.

Received: 7 June 2023 Accepted: 29 November 2023

Published online: 08 December 2023

References

- Ritchie MD. Large-scale analysis of genetic and clinical patient data. *Ann Rev Biomed Data Sci.* 2018;1(1):263–74.
- Stephens ZD, et al. Big data: astronomical or genetical? *PLoS Biol.* 2015;13(7):e1002195.
- Zhu H. Big data and artificial Intelligence modeling for drug discovery. *Annu Rev Pharmacol Toxicol.* 2020;60(1):573–89.
- Wilkinson MD, Dumontier M, et al. The FAIR Guiding Principles for scientific data management and stewardship. *Sci Data.* 2016;3(1):160018.
- Waagmeester A, Stupp G, et al. Wikidata as a knowledge graph for the life sciences. *eLife.* 2020;9:e52614.
- Belleau F, Nolin MA, Tourigny N, Rigault P, Morissette J. Bio2RDF: towards a mashup to build bioinformatics knowledge systems. *J Biomed Inform.* 2008;41(5):706–16.
- Williams AJ, Harland L, Groth P, Pettifer S, Chichester C, Willighagen EL, et al. Open PHACTS: semantic interoperability for drug discovery. *Drug Discov Today.* 2012;17(21–22):1188–98.
- Domingo-Fernández D, Baksi S, et al. COVID-19 Knowledge Graph: a computable, multi-modal, cause-and-effect knowledge model of COVID-19 pathophysiology. *Bioinformatics.* 2020. p. 1332–4.
- Himmelstein DS, Lizee A, et al. Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *eLife.* 2017;6:e26726.
- Hogan A, Blomqvist E, Cochez M, d'Amato C, de Melo G, Gutiérrez C, et al. Knowledge Graphs. No. 22 in *Synthesis Lectures on Data, Semantics, and*

- Knowledge. Springer; 2021. <https://doi.org/10.2200/S01125ED1V01Y202109DSK022>. <https://kgbook.org/>.
11. Chichester C, Digles D, et al. Drug discovery FAQs: workflows for answering multidomain drug discovery questions. *Drug Discov Today*. 2015;20(4):399–405.
 12. Knox C, Law V, Jewison T, Liu P, et al. DrugBank 3.0: a Comprehensive Resource for omics' Research on Drugs. *Nucleic Acids Res*. 2010;39(suppl_1):D1035–41.
 13. Bateman A, Martin MJ, O'Donovan C, Magrane M, Apweiler R, Alpi E, et al. UniProt: a hub for protein information. *Nucleic Acids Res*. 2015;43(Database issue):D204–12.
 14. Lipscomb CE. Medical Subject Headings (MeSH). *Bull Med Libr Assoc*. 2000;88(3): 265–6.
 15. Morselli Gysi D, Do Valle Í, Zitnik M, Ameli A, Gan X, Varol O, et al. Network medicine framework for identifying drug-repurposing opportunities for COVID-19. *Proc Natl Acad Sci*. 2021;118(19):e2025581118.
 16. Nickel M, Murphy K, Tresp V, Gabrilovich E. A Review of Relational Machine Learning for Knowledge Graphs. *Proc IEEE*. 2016;104(1):11–33. <https://doi.org/10.1109/JPROC.2015.2483592>.
 17. Wang Q, Mao Z, Wang B, Guo L. Knowledge Graph Embedding: A Survey of Approaches and Applications. *IEEE Trans Knowl Data Eng*. 2017;29(12):2724–43. <https://doi.org/10.1109/TKDE.2017.2754499>.
 18. Bordes A, Usunier N, García-Durán A, Weston J, Yakhnenko O. Translating Embeddings for Modeling Multi-relational Data. In: Burges CJC, Bottou L, Ghahramani Z, Weinberger KQ, editors. *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5–8, 2013, Lake Tahoe, Nevada, United States; 2013*. p. 2787–2795. <https://proceedings.neurips.cc/paper/2013/hash/1cecc7a77928ca8133fa24680a88d2f9-Abstract.html>.
 19. Trouillon T, Welbl J, Riedel S, Gaussier É, Bouchard G. Complex Embeddings for Simple Link Prediction. In: Balcan M, Weinberger KQ, editors. *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19–24, 2016*. vol. 48 of *JMLR Workshop and Conference Proceedings*. JMLR.org; 2016. p. 2071–2080. <http://proceedings.mlr.press/v48/trouillon16.html>.
 20. Sun Z, Deng Z, Nie J, Tang J, RotatÉ. Knowledge Graph Embedding by Relational Rotation in Complex Space. In: *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*. OpenReview.net; 2019. <https://openreview.net/forum?id=HkgEQnRqYQ>.
 21. Xie R, Liu Z, Jia J, Luan H, Sun M. Representation Learning of Knowledge Graphs with Entity Descriptions. In: Schuurmans D, Wellman MP, editors. *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12–17, 2016, Phoenix, Arizona, USA*. AAAI Press; 2016. p. 2659–2665. <http://www.aaai.org/ocs/index.php/AAAI/AAAI16/paper/view/12216>.
 22. Teru KK, Denis EG, Hamilton WL. Inductive Relation Prediction by Sub-graph Reasoning. In: *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event*. vol. 119 of *Proceedings of Machine Learning Research*. PMLR; 2020. p. 9448–9457. <http://proceedings.mlr.press/v119/teru20a.html>.
 23. Daza D, Cochez M, Groth P. Inductive Entity Representations from Text via Link Prediction. In: Leskovec J, Grobelnik M, Najork M, Tang J, Zia L, editors. *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19–23, 2021*. ACM / IW3C2; 2021. p. 798–808. <https://doi.org/10.1145/3442381.3450141>.
 24. Galkin M, Denis EG, Wu J, Hamilton WL. NodePiece: Compositional and Parameter-Efficient Representations of Large Knowledge Graphs. In: *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25–29, 2022*. OpenReview.net; 2022. p. 1–14. <https://openreview.net/forum?id=xMJWUKJnFSw>.
 25. Xie R, Liu Z, Luan H, Sun M. Image-embodied Knowledge Representation Learning. In: Sierra C, editor. *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19–25, 2017*. ijcai.org; 2017. p. 3140–3146. <https://doi.org/10.24963/ijcai.2017/438>.
 26. Tay Y, Tuan LA, Phan MC, Hui SC. Multi-Task Neural Network for Non-discrete Attribute Prediction in Knowledge Graphs. In: Lim E, Winslett M, Sanderson M, Fu AW, Sun J, Culpepper JS, et al., editors. *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, CIKM 2017, Singapore, November 06 – 10, 2017*. ACM; 2017. p. 1029–1038. <https://doi.org/10.1145/3132847.3132937>.
 27. Wu Y, Wang Z. Knowledge Graph Embedding with Numeric Attributes of Entities. In: *Proceedings of the Third Workshop on Representation Learning for NLP, Melbourne, Australia: Association for Computational Linguistics; 2018*. p. 132–136. <https://doi.org/10.18653/v1/W18-3017>. <https://aclanthology.org/W18-3017>.
 28. Pezeshkpour P, Chen L, Singh S. Embedding Multimodal Relational Data for Knowledge Base Completion. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics; 2018. p. 3208–3218. <https://doi.org/10.18653/v1/D18-1359>. <https://aclanthology.org/D18-1359>.
 29. Kristiadi A, Khan MA, Lukovnikov D, Lehmann J, Fischer A. Incorporating Literals into Knowledge Graph Embeddings. In: Ghidini C, Hartig O, Maleshkova M, Svátek V, Cruz IF, Hogan A, et al., editors. *The Semantic Web - ISWC 2019 - 18th International Semantic Web Conference, Auckland, New Zealand, October 26–30, 2019, Proceedings, Part I*. vol. 11778 of *Lecture Notes in Computer Science*. Springer; 2019. p. 347–363. https://doi.org/10.1007/978-3-030-30793-6_20.
 30. Wang X, Gao T, Zhu Z, Zhang Z, Liu Z, Li J, et al. KEPLER: A Unified Model for Knowledge Embedding and Pre-trained Language Representation. *Trans Assoc Comput Linguistics*. 2021;9:176–94. https://doi.org/10.1162/tacl_a_00360.
 31. Ektefaie Y, Dasoulas G, Noori A, Farhat M, Zitnik M. Multimodal learning with graphs. *Nat Mach Intel*. 2023. <https://doi.org/10.1038/s42256-023-00624-6>.
 32. Wang L, Zhao W, Wei Z, Liu J. SimKGC: Simple Contrastive Knowledge Graph Completion with Pre-trained Language Models. In: Muresan S, Nakov P, Villavicencio A, editors. *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22–27, 2022*. Association for Computational Linguistics; 2022. p. 4281–4294. <https://doi.org/10.18653/v1/2022.acl-long.295>.
 33. Markowitz E, Balasubramanian K, Mirtaheer M, Annavaram M, Galstyan A, Steeg GV. STAIK: Structure and Text for Inductive Knowledge Graph Completion. In: Carpuat M, de Marneffe M, Ruiz IVM, editors. *Findings of the Association for Computational Linguistics: NAACL 2022, Seattle, WA, United States, July 10–15, 2022*. Association for Computational Linguistics; 2022. p. 604–615. <https://doi.org/10.18653/v1/2022.findings-naacl.46>.
 34. Safavi T, Downey D, Hope T. CascadER: Cross-Modal Cascading for Knowledge Graph Link Prediction. *CoRR*. 2022. <https://doi.org/10.48550/arXiv.2205.08012>.
 35. Devlin J, Chang M, Lee K, Toutanova K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In: Burstein J, Doran C, Solorio T, editors. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2–7, 2019, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics; 2019. p. 4171–4186. <https://doi.org/10.18653/v1/n19-1423>.
 36. Ali M, Hoyt CT, ndez D, Lehmann J, Jabeen H. BioKEEN: a library for learning and evaluating biological knowledge graph embeddings. *Bioinformatics*. 2019;35(18):3538–40.
 37. Nelson W, Zitnik M, Wang B, Leskovec J, Goldenberg A, Sharan R. To embed or not: network embedding as a paradigm in computational biology. *Front Genet*. 2019;10:381.
 38. Walsh B, Mohamed SK, Nováček V. BioKG: A knowledge graph for relational learning on biological data. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. New York: Association for Computing Machinery; 2020. p. 3173–80.
 39. Mohamed SK, ek V, Nounu A. Discovering protein drug targets using knowledge graph embeddings. *Bioinformatics*. 2020;36(2):603–10.
 40. Alshahrani M, Thafar MA, Essack M. Application and evaluation of knowledge graph embeddings in biomedical data. *PeerJ Comput Sci*. 2021;7:e341.
 41. Ye C, Swiers R, Bonner S, Barrett I. A Knowledge Graph-Enhanced Tensor Factorisation Model for Discovering Drug Targets. *IEEE/ACM Trans Comput Biol Bioinform*. 2022. p. 3070–80.

42. Gema AP, Grabarczyk D, De Wulf W, Borole P, Alfaro JA, Minervini P, et al. Knowledge Graph Embeddings in the Biomedical Domain: Are They Useful? A Look at Link Prediction, Rule Learning, and Downstream Polyp-harmacy Tasks. *CoRR*. 2022. <https://doi.org/10.48550/arXiv.2305.19979>.
43. Karim MR, Cochez M, Jares JB, Uddin M, Beyan OD, Decker S. Drug-Drug Interaction Prediction Based on Knowledge Graph Embeddings and Convolutional-LSTM Network. In: Shi XM, Buck M, Ma J, Veltri P, editors. Proceedings of the 10th ACM International Conference on Bioinformatics, Computational Biology and Health Informatics, BCB 2019, Niagara Falls, NY, USA, September 7-10, 2019. ACM; 2019. p. 113–123. <https://doi.org/10.1145/3307339.3342161>.
44. Choi W, Lee H. Identifying disease-gene associations using a convolutional neural network-based model by embedding a biological knowledge graph with entity descriptions. *PLoS ONE*. 2021;16(10):e0258626.
45. Alshahrani M, Almansour A, Alkhalidi A, Thafar MA, Uludag M, Essack M, et al. Combining biomedical knowledge graphs and text to improve predictions for drug-target interactions and drug-indications. *PeerJ*. 2022;10:e13061.
46. Ren Z, You Z, Yu C, Li L, Guan Y, Guo L, et al. A biomedical knowledge graph-based method for drug-drug interactions prediction through combining local and global features with deep neural networks. *Briefings Bioinform*. 2022;23(5). <https://doi.org/10.1093/bib/bbac363>.
47. Su X, Hu L, You Z, Hu P, Zhao B. Attention-based knowledge graph representation learning for predicting drug-drug interactions. *Brief Bioinform*. 2022;23(3):bbac140.
48. Zhang Y, Li Z, Duan B, Qin L, Peng J. MKGE: Knowledge graph embedding with molecular structure information. *Comput Biol Chem*. 2022;100:107730. <https://doi.org/10.1016/j.compbiolchem.2022.107730>.
49. Zhu C, Yang Z, Xia X, Li N, Zhong F, Liu L. Multimodal reasoning based on knowledge graph embedding for specific diseases. *Bioinform*. 2022;38(8):2235–45. <https://doi.org/10.1093/bioinformatics/btac085>.
50. Ruffinelli D, Broscheit S, Gemulla R. You CAN Teach an Old Dog New Tricks! On Training Knowledge Graph Embeddings. In: 8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020. OpenReview.net; 2020. <https://openreview.net/forum?id=BkxSmlBFvr>.
51. Ali M, Berrendorf M, Hoyt CT, Vermue L, Galkin M, Sharifzadeh S, et al. Bringing Light Into the Dark: A Large-Scale Evaluation of Knowledge Graph Embedding Models Under a Unified Framework. *IEEE Trans Pattern Anal Mach Intell*. 2022;44(12):8825–45. <https://doi.org/10.1109/TPAMI.2021.3124805>.
52. Bonner S, Barrett IP, Ye C, Swiers R, Engkvist O, Hoyt CT, et al. Understanding the performance of knowledge graph embeddings in drug discovery. *Artif Intell Life Sci*. 2022;2:100036. <https://doi.org/10.1016/j.aillsci.2022.100036>. <https://www.sciencedirect.com/science/article/pii/S266731852000071>
53. Wishart DS, Knox C, Guo AC, Cheng D, Shrivastava S, Tzur D, et al. DrugBank: a knowledgebase for drugs, drug actions and drug targets. *Nucleic Acids Res*. 2008;36(Database issue):D901–6.
54. Yamanishi Y, Araki M, Gutteridge A, Honda W, Kanehisa M. Prediction of drug-target interaction networks from the integration of chemical and genomic spaces. *Bioinformatics*. 2008;24(13):i232–40.
55. Elnaggar A, Heinzinger M, Dallago C, Rehawi G, Wang Y, Jones L, et al. Prottrans: toward understanding the language of life through self-supervised learning. *IEEE Trans Pattern Anal Mach Intell*. 2021;44(10):7112–27.
56. Morris P, St Clair R, Hahn WE, Barenholtz E. Predicting Binding from Screening Assays with Transformer Network Embeddings. *J Chem Inf Model*. 2020. <https://doi.org/10.1021/acs.jcim.9b01212>.
57. Lee J, Yoon W, Kim S, Kim D, Kim S, So CH, et al. BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*. 2019;36(4):1234–40. <https://doi.org/10.1093/bioinformatics/btz682>.
58. Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is All you Need. In: Guyon I, von Luxburg U, Bengio S, Wallach HM, Fergus R, Vishwanathan SVN, et al., editors. Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4–9, 2017, Long Beach, CA, USA; 2017. p. 5998–6008. <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.
59. Jumper J, Evans R, Pritzel A, Green T, Figurnov M, Ronneberger O, et al. Highly accurate protein structure prediction with AlphaFold. *Nature*. 2021;596(7873):583–9.
60. Ali M, Berrendorf M, Hoyt CT, Vermue L, Sharifzadeh S, Tresp V, et al. PyKEEN 1.0: A Python Library for Training and Evaluating Knowledge Graph Embeddings. *J Mach Learn Res*. 2021;22:82:1–6. <http://jmlr.org/papers/v22/20-825.html>.
61. Ali M, Matinata A. Knowledge graph embeddings: Are relation-learning models learning relations? In: EDBT/ICDT Workshops. Aachen: CEUR-WS.org; 2020.
62. Nascimento AC, Prudêncio RB, Costa IG. A multiple kernel learning algorithm for drug-target interaction prediction. *BMC Bioinformatics*. 2016;17:1–16.
63. Hao M, Bryant SH, Wang Y. Predicting drug-target interactions by dual-network integrated logistic matrix factorization. *Sci Rep*. 2017;7(1):1–11.
64. Olayan RS, Ashoor H, Bajic VB. DDR: efficient computational method to predict drug-target interactions using graph mining and machine learning approaches. *Bioinformatics*. 2018;34(7):1164–73.
65. Takaya MS, Rehmsmeier. The Precision-Recall Plot Is More Informative than the ROC Plot When Evaluating Binary Classifiers on Imbalanced Datasets. *PLoS ONE*. 2015;10:1–21. <https://doi.org/10.1371/journal.pone.0118432>.
66. Bonner S, Kirik U, Engkvist O, Tang J, Barrett IP. Implications of topological imbalance for representation learning on biomedical knowledge graphs. *Brief Bioinform*. 2022;23(5). Bbac279. <https://doi.org/10.1093/bib/bbac279>.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Ready to submit your research? Choose BMC and benefit from:

- fast, convenient online submission
- thorough peer review by experienced researchers in your field
- rapid publication on acceptance
- support for research data, including large and complex data types
- gold Open Access which fosters wider collaboration and increased citations
- maximum visibility for your research: over 100M website views per year

At BMC, research is always in progress.

Learn more biomedcentral.com/submissions

