



A Multi-modal Knowledge Graph Platform Based on Medical Data Lake

Ruoyu Wang^(✉)

School of Methematical Sciences, Beihang University, Beijing, China
18377446@buaa.edu.cn

Abstract. In recent years, the application of Knowledge Graphs (KGs) in the medical domain has developed rapidly. Embedding medical knowledge graphs into low-dimensional vectors to implement certain applications is the focus of recent research. However, current approaches focus on using structured data while ignoring external information from prior knowledge and multi-modal data, such as visual and textual data. In this paper, we propose a multi-modal knowledge graph platform based on medical data lake. It enhances Knowledge Graph Embedding (KGE) from structural data by using prior knowledge like doctors' advice. After the generation of embedding of knowledge graphs, we use a fully-connected layer to combine the structural vectors and multi-modal embedding from unstructured data for link prediction. Moreover, we provide recommendations to help doctors diagnose and research based on the knowledge graph embedding model. Our platform has been used in a clinical decision support system, which is capable of discovering unveiled relations between entities and giving suitable advice for doctors.

Keywords: Knowledge graph · Medical data lake · Multi-modal data

1 Introduction

Over the past few years, research on knowledge graph (KG) has thrived, which produced a significant number of remarkable achievements. A wide variety of applications are implemented, including recommendations system, Q&A and information retrieval [2, 8]. Knowledge graph is a type of knowledge base that uses a graph-structured data model to store semantic information in the form of entities and the relationship between them. It is composed of many triplets (head entity-relation-tail entity) and each triplet represents one piece of knowledge.

In the medical domain, there are different types of entities, such as diseases, symptoms, treatments, etc and a variety of relationships between them. The information in the knowledge graph can be supplemented by prior knowledge from some physicians and data from other models, which assists providers by validating diagnoses and identifying treatment plans based on individual needs. When it comes to applications like Q&A and recommendation systems, the key point is that the accuracy of the answers and recommendations is tightly linked

with the correctness and completeness of the knowledge graphs. Most of the popular knowledge graphs (e.g., YAGO, DBPedia, or Wikidata) remain incomplete despite the great effort invested in their creation and maintenance, even though they have already stored large amounts of entities and relationships. To complete the knowledge graph, researchers proposed the knowledge graph embedding (KGE) method, which is to embed the entities and relations into low-dimensional vectors. In the KGE field, the most representative translational distance model is TransE [1]. Many researchers have proposed a significant number of improved models to increase its accuracy, such as TransH [17] and TransD [6]. Apart from that, researchers have also tried to use the information from multi-modal data to supplement the embedding [12]. They have tried to integrate textual information [19] and visual information [20] into knowledge graphs. However, it is rare to include multi-modal data in existing medical knowledge graphs.

With the advances of medical technology, enormous amounts of multi-modal medical data are generated, and traditional database is not able to store and manage the data efficiently. A new type of database called data lake is good at handling this problem [14, 21]. In data lake, various types of data can be stored and processed, such as Electronic Medical Records (EMRs), Magnetic Resonance Imaging (MRI) scans, Computerized Tomography (CT) scans, X-ray and Positron Emission Tomography (PET) scans, etc. Moreover, doctors can give their advice and document their experience in it. All of the metadata can be saved for traceability and safety in data lake. For example, images like magnetic resonance imaging (MRI) can be embedded into embedding vectors, and the vectors will be stored to represent specific features. Due to the powerful storage capacity and convenient operations, data lake supports plenty of applications like data exploration. Since its advantages in data exploration and abundant multi-modal data, we can efficiently construct a knowledge graph based on it, and import the multi-modal data and documented experiences to improve its quality. Therefore, we propose a Multi-Modal Knowledge Graph Platform (MMKGP) based on medical data lake.

Then this knowledge graph can provide advice and recommendations for doctors. We address our contributions as follows:

1. A multi-modal knowledge graph platform that uses structured data, multi-modal data and doctors' experiences to construct a highly complete and accurate knowledge graph based on medical data lake.
2. A translation-based model of KGE which can make use of the constraints of pairs of relations according to prior knowledge.
3. A method to use external information in multi-modal medical data to find missing relations of knowledge graph.
4. The platform has been used in a clinical decision support system, which is capable of discovering unveiled relations between entities and giving suitable advice for doctors

In the following, we will first review the related work in Sect. 2, and then introduce the overview of our proposed framework in Sect. 3. How to use prior knowledge to get more accurate embedding is demonstrated in Sect. 4, and how

to find missing relations with the help of unstructured data is presented in Sect. 5. After that, we will explain how to use the generated high-quality knowledge graph to help doctors in Sect. 6 followed by the conclusion in Sect. 7.

2 Related Work

In this chapter, we discuss the knowledge graph embedding and data lake.

There are many knowledge graph embedding models. In 2013, TransE model was proposed [1], which considers relations as translating operations between head and tail entities. A lot of translation-based models were proposed to achieve a better performance in representation learning tasks of knowledge graph after TransE, such as link prediction [9]. For instance, TransH [17], which considers relations as the vectors on a hyperplane, was proposed to overcome TransE’s drawbacks in dealing with 1-to-N, N-to-1 and N-to-N relations. However, entities and relations are still in the same semantic space, which limits the modeling ability. TransR [9] appeared and assumed that entities and relationships belong to different semantic spaces. Furthermore, TransD [6] was proposed to simplify TransR by projection matrix. Besides translation-based models, semantic matching models (such as RESCAL [13], where the entire knowledge graph is encoded as a 3D tensor) and graph neural networks models (such as R-GCN [15], aggregating different nodes according to different relationships) are other good ways to represent knowledge graphs.

Later attempts focus on using background knowledge to constraint the representation learning [4, 5, 16] or integrating extra information beyond triplets [12]. In 2018, Ding et al. noticed the connection between the pairs of relations [4], and proposed proposes to use the background information to constrain the embeddings of relations in RESCAL. Xie et al. focused on using the extra information from multi-modal data to supplement the embedding of knowledge graphs. They proposed methods to make use of textual description [19] and image information [20] in the construction of knowledge graphs. In the medical domain, myDIG [7] and SemTK [3] proposed knowledge graph building tools which can extract triplets from website and texts. EMKN [22] constructed an EMR-based medical knowledge graph by extracting the medical entities. Liu et al. proposed to extract triplets from 1454 clinically pediatric cases, and then combined them with expert experiences and book knowledge in the generation of knowledge graphs [10].

However, it is not enough to combine the translation-based model just with the multi-modal information or just with prior knowledge. There is no model that combines structured, multi-modal information and prior knowledge together to build a medical knowledge graph (See Table 1).

Last decade has witnessed great development of the application of knowledge graphs. Bordes et al. proposed to embed knowledge graphs into vectors and use these vectors to do link prediction which can specify the unrevealed relations between the entities [1]. Apart from that, we can also break through the boundary of the input knowledge graph and expand the application into

Table 1. Comparison of the information contained in the model

| Method | Triples | Multi-modal | Prior knowledge |
|------------|---------|-------------|-----------------|
| TransE [1] | + | - | - |
| IKRL [20] | + | + | - |
| CNA [4] | + | - | + |
| myDIG [7] | + | - | - |
| SemTK [3] | + | - | - |
| EMKN [22] | + | - | - |
| HKDP [10] | + | - | + |
| Ours | + | + | + |

Note: (+) means this kind of information is contained in the model and (-) means this type of this kind of information is not contained in the model.

boarder domains. In 2013, Weston et al. proposed to use knowledge graphs to extract unseen relation between two entities [18]. Another important application of knowledge graph is the wide use in the recommender systems [8]. We can use knowledge graph as a big database and give recommendation for the users according to the their preference and history operation. Moreover, knowledge graph can be also used in the Q&A systems [2], which is very helpful in the medical domain.

To store and manage multi-modal data in the medical domain, Zhang et al. proposed to use medical data lake to deal with unstructured data [21]. As for unstructured data, data lake embeds it into vectors by state-of-the-art models automatically, and then it can be regarded as a sub-tree that contains the nodes representing its path information, feature vectors, and others. Data lake provides a convenient way for us to query data and explore data efficiently in the construction of knowledge graph.

3 Architecture of MMKGP

MMKGP is established on the medical data lake, in which three sources of information are available: prior knowledge provided by doctors, structured data or primary knowledge graph, unstructured data such as linguistic and visual information. The architecture of MMKGP is presented in Fig. 1.

The experience from medical experts can be understood as prior knowledge, and we can learn the constraints of pairs of relations and their weights from it. The structured data is put into the TransE model, then we can get the structured embedding of entities and relations. After that, the structured embedding and the weights of pairs of constrained relations are imported into the first model to generate the updated structured embedding. Next, based on the unstructured embedding, we construct a fully-connected layer to complement the missing relations between pairs of entities. Finally, a highly-quality knowledge graph can be

generated, which includes newly found relations, and can help doctors make diagnoses and research.




Fig. 1. Overview

4 Translation-Based Model Enhanced by Prior Knowledge

We take doctors’ professional experience as prior knowledge in the construction of knowledge graph. At the same time we mine the constraints of the relations from the knowledge graph automatically [4]. For the two types of prior knowledge, we use the same method to represent it, which is the constraint weight of pairs of relations. The weight of relation A and relation B represents the certainty that we can infer B from A. For example, if knowledge graph has triplet (h, r_1, t) and the constraint weight of r_1 and r_2 is -0.95 , we can conclude that triplet (h, r_2, t) is correct with 95% certainty.

In the traditional translation-based model, the constraints of relations is usually ignored. In fact, the constraints can provide a lot of information. For example, there are two diseases called diabetes mellitus and coronary heart disease, and we know that diabetes mellitus is “cause of” Mike (See Fig. 2). So we have one triplet $(diabetes_mellitus, cause_of, coronary_heart_disease)$ in the knowledge graph. Due to the relationship between “cause_of” and “complication_of”, we need to control the embedding vector of “complication_of”

is nearly the reverse of the embedding vector of “cause_of”. And then, based on the principle of TransE, we can easily conclude that coronary heart disease is “complication_of” diabetes mellitus so that we can add up a new triplet (*coronary_heart_disease, complication_of, diabetes_mellitus*) into the knowledge graph. It is just a simple example, but we can understand the importance of constraints in KGE and why these constraints can significantly improve the accuracy in link prediction.

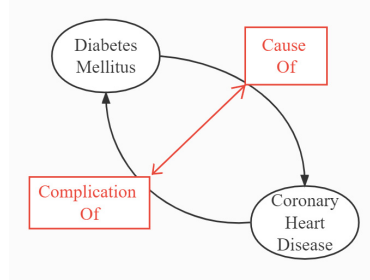


Fig. 2. Constrained relations

4.1 TransE Model

In traditional translation-based model, given an existing triplet (h, r, t) where h represents the head entity, r represents the relation and t represents the tail entity, the relation is embedded into a translation vector \mathbf{r} so that the embedded head entity \mathbf{h} and the embedded tail entity \mathbf{t} can be connected by \mathbf{r} in low error, e.g., $\mathbf{h} + \mathbf{r} \approx \mathbf{t}$. For example, if triplet $(COVID-19, type_of, pneumonia)$ holds, after the embedding in model TransE, we are looking forward to the following equation satisfied in low error:

$$\mathbf{V}_{COVID-19} + \mathbf{V}_{type_of} \approx \mathbf{V}_{pneumonia},$$

where $V_{r/e}$ means the embedding vector of the relation r or the entity e . The scoring function of TransE is defined as the distance between $\mathbf{h} + \mathbf{r}$ and \mathbf{t} :

$$f_r(h, t) = \|\mathbf{h} + \mathbf{r} - \mathbf{t}\|_{\frac{1}{2}},$$

The score is expected to be small if the triplet (h, r, t) is correct.

4.2 Constraint for Relations

Based on TransE, we propose to add the constraints for relations. In the medical domain, on the one hand, we can get the entailment of relations from doctors’ experience. On the other hand, the constraints can be automatically calculated

by modern rule mining system [11]. We use a weight to represent the degree that the former approximately constraint the latter, e.g., PARENT_OF and CHILD_OF. The maximum of the absolute value of the weight is 1, and the minimum is 0. The larger the value is, the more restrictive the latter is by the former and the more similar this pair of the relations' embedding vectors should be. The sign means the direction of the latter relation. If the sign is negative, we should calculate the constraint of the reversed latter vector with the former vector.

We first explore the strict constraint, e.g., weight $\lambda = 1$. This strict constraint $r_1 \rightarrow r_2$ means that if relation r_1 holds, then relation r_2 holds. And the constraint can be roughly represented in the following equation:

$$\mathbf{r}_1 = \mathbf{r}_2$$

where \mathbf{r}_i means the embedding vector of the relation r_i .

For non-strict constraints, e.g., weight $\lambda = 0.95$, we use a power function to measure the entailment of the vectors:

$$f_\lambda(\mathbf{r}_1, \mathbf{r}_2) = \|\mathbf{r}_1 - \mathbf{r}_2\|_{\frac{1}{2}}^{\|\lambda\|},$$

where \mathbf{r}_i means the embedding vector of the relation r_i and λ is the value of constraint weight. When weight $\lambda = \pm 1$, which means the constraint is strict, this function degenerated into the distance between \mathbf{r}_1 and \mathbf{r}_2 . Due to the extension to the model, we should also add up the scoring function of constrained relations. The scoring function is defined in the following equation:

$$[f_\lambda(\mathbf{r}_1, \mathbf{r}_2) = \|\mathbf{r}_1 - \mathbf{r}_2\|_{\frac{1}{2}}^{\|\lambda\|} - \gamma]_+,$$

where γ is the margin value can be given by users and $[x]_+ = \max(0, x)$. During the training, we not only calculate the scoring function of triplets, but also calculate the scoring function of constrained relations, and add them up into one combined scoring function. And then, we use the stochastic gradient descent (SGD) to adjust parameters to lower the function.

5 Knowledge Graph Completion with Multi-modal Data

Our method is based on the implementation of data lake. The data in the data Lake includes structured data, semi-structured data, unstructured data and binary data, so as to form a centralized data storage containing all forms of data. We add a fully connected layer into this model so that it can consider the external information from entities when the model needs to make link prediction or triplet classification. There is much external information in the multi-modal data which can help us improve our knowledge graph embedding. Picture or linguistic information stored in data lake can correct the wrong relationship or add some new relationships.

For example, a person who suffers from influenza has symptoms of cough and runny nose (See Fig. 3). The triplet is (cough, is_symptom_of, influenza)

and (runny_nose, is_symptom_of, influenza). A coughing person needs cough syrup (cough, need_medicine, cough_syrup). TransE may predict that a person with a runny nose only will also need cough syrup (runny_nose, need_medicine, cough_syrup). But we get a picture of a prescription showing that doctors use nasal spray instead of cough syrup to treat runny nose patients. After the visual embedding, we can amend the relation among runny nose, cough syrup and nasal spray (runny_nose, need_medicine, nasal_spray).

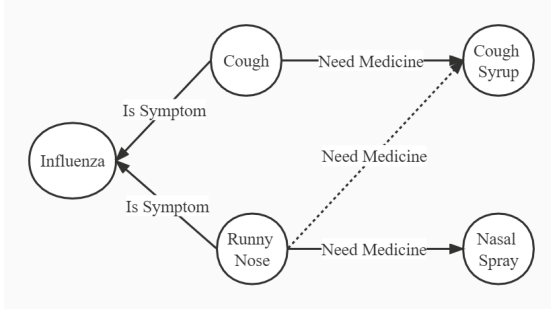


Fig. 3. An example using multi-modal embedding.

5.1 Dataset

Our data is collected from three hospitals, including more than 500 diseases and medical records of 3 million patients. This dataset includes 8293284 concepts, 83591932 entities and 295848293 relationships, of which 32256360 are the relationships between concepts.

5.2 Evaluation Criterion

Triple Energy: The energy includes three parts:

- Structural Energy: we set the energy function in terms of the rules of TransE as $E_s = \|h_s + r_s - t_s\|$.
- Multi-modal Energies: First, we define the multi-modal representations h_m and t_m of the head and the tail entities. There are two parts of multi-modal energy functions: $E_{m1} = \|h_m + r_s - t_m\|$ and $E_{m2} = \|(h_m + h_s) + r_s - (t_m + t_s)\|$.
- Structural-multi-modal Energies: We need to make sure structural and the multi-modal representations are in the same space, so we define the energy functions as: $E_{sm} = \|h_s + r_s - t_m\|$ and $E_{ms} = \|h_m + r_s - t_s\|$.

Finally, we add all of them to define our triple energy.

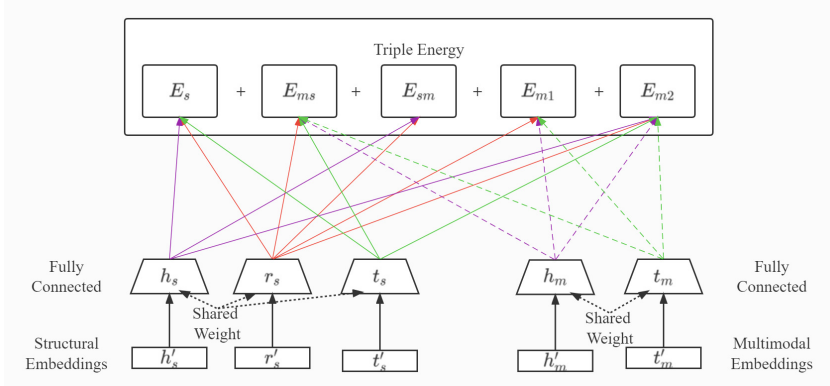


Fig. 4. Overview of the fully connect layer and triple energy

Objective Function: First, we need two negative triple sets. One is to replace head entities named T'_{head} , and another is to replace the tail entities named T'_{tail} . Then we set $L = L_h + L_t$ to represent the loss.

$$L_h = \sum_T \sum_{T'_{tail}} MAX(E(h, r, t) - E'(h, r, t'))$$

$$L_t = \sum_T \sum_{T'_{head}} MAX(E(h, r, t) - E'(h', r, t))$$

Finally, in order to consider various information during knowledge graph embedding, we set fully connected layer that can map both structural and multi-modal representations into the same place (See Fig. 4). Weights are also shared between those inputs.

5.3 Model Training and Result

Model Training: After completing these preparations, we can start the model training.

Our method is divided into two parts. The first is to train TransE+AER. We set batch size to 100, embedded size to 100 dimensions, learning rate to 1.0 and use PyTorch to carry out the random gradient descent method. With the help of OpenKE, we take part of the positive and negative triples in each iteration. Then update the parameters to enter the next iteration. The loss of each iteration is displayed on the terminal panel. The model saves the real-time model parameters and optimal model parameters. Before the next part training, we test this embedding result to make a control group.

In the second part, we use the structural embeddings to continue our training. We set batch size to 100, initial learning rate to 0.001, and use the Adam optimizer to get the best results. We combine linguistic embeddings and visual embeddings to get the final KG.

Table 2. Link prediction results

| Method | Raw | | Filter | |
|--------------|-------|---------|--------|---------|
| | MR | Hits@10 | MR | Hits@10 |
| TransE | 160.0 | 0.641 | 152.1 | 0.643 |
| TransE+AER | 148.3 | 0.720 | 145.6 | 0.720 |
| TransE+Multi | 25.2 | 0.666 | 25.02 | 0.672 |
| Ours | 18.7 | 0.766 | 18.6 | 0.772 |

Comparison of Various Methods: We compared our result with several methods, including TransE, TransE+AER, TransE+multi-modal. Table 2 shows the results. Our idea combining the advantages of both those methods can do better in knowledge graph embedding and Link Prediction. We get a lower MR and higher Hits@10. Prior knowledge helps us add more information of relations so that knowledge graph becomes more comprehensive. Multi-modal data can help correct the unsuitable relations between entities. Therefore, we get a more detailed and accurate knowledge graph.

6 Knowledge Graph-Based Clinical Decision Support System

To help doctors to diagnose and do research, we have generated a knowledge graph and tried to improve its quality by using our platform. During the quality improvement, new triplets considering as pieces of knowledge will be generated. They can clarify the relations between pairs of specific entities. The newly derived knowledge can help doctors to find new connections between entities. After the improvement, we get a high-quality knowledge graph which can efficiently and accurately perform tasks like recommendation and Q&A. According to the above, we can divide the applications of the knowledge graph into applications during improvement and applications after improvement.

6.1 Link Prediction & Correction

While constructing a high-quality knowledge graph, we have done link prediction to complete the knowledge graph, which can help doctors find out the missing relationship between two entities. The system stores these newly derived triplets into a list to save them. After the completion, this list will be printed out and shown to the doctors and researchers. They can check it carefully and try to point out the potential reason for these pieces of knowledge, which may give researchers clues and ideas.

Apart from the completion, this system can also check the accuracy of existing knowledge. Doctors can put forward suggestions to modify the existing triples according to their own knowledge. After checking every triplet, the list will be printed out to the doctors so that they can check whether any mistakes happen

or the data is wrong. Then the staff or doctors can correct the wrong data to reduce potential risks. Some existing knowledge may be wrong. For example, this system may hint to doctors that some of the treatments for a specific disease are inappropriate. Especially if doctors have corresponding conjecture, this hint may lead to a retest to the treatment and may help patients get proper curing.

6.2 Recommendation and Q&A System

One application is to use this system to recommend treatments or medicines according to the patients' basic information like age, sex, weight, etc. and their symptoms. For example, in the meeting with a patient, we import his name and symptoms. After checking such as CT, this system can extract the information from the CT image. By comprehensive analysis of symptoms and the checking reports, this system can clearly understand the condition of the patient. Then it can give diagnosis to the disease suffered by the patient and its recommendations to the corresponding treatment. The system also considers the age and the weight of the patient and gives its recommendation to the dose of medicine.

Apart from recommendation, knowledge graph supports Q&A system [2] as well. During patient meetings, diagnose and research, doctors can type their questions into this system. Then state-of-the-art models are used to get the potential answer of the question on the knowledge graph. This is especially helpful for the new graduate or intern doctors because they don't have many experiences. It can also be applied to the online consultation services. People with some simple symptoms can query online and get the answer quickly so that they don't need to spend a lot of time going to a hospital and wait for the meeting with doctors.

7 Conclusion

We present a multi-modal platform to construct high-quality knowledge graph-based on medical data lake. Based on the fact that experts' experience can give essential guidance in the medical domain, we propose to make use of this experience to calculate the constraint weights for the pairs of relation. These weights are used to enhance the embedding vector of the relations. Considering the neglected external information from unstructured data, we construct a fully-connected layer to combine both structured and unstructured embedding of data to find the missing relations between entities.

Acknowledgement. This work was supported by National Key R&D Program of China(2020AAA0109603).

References

1. Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., Yakhnenko, O.: Translating embeddings for modeling multi-relational data. In: Burges, C.J.C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 26. Curran Associates, Inc. (2013)

2. Bordes, A., Weston, J., Usunier, N.: Open Question Answering with Weakly Supervised Embedding Models. Springer, New York, Inc (2014)
3. Cuddihy, P., Mchugh, J., Williams, J.W., Mulwad, V., Aggour, K.S.: Semtk: an ontology-first, open source semantic toolkit for managing and querying knowledge graphs (2017)
4. Ding, B., Wang, Q., Wang, B., Guo, L.: Improving knowledge graph embedding using simple constraints. In: Meeting of the Association for Computational Linguistics (2018)
5. Guo, S., Wang, Q., Wang, L., Wang, B., Guo, L.: Jointly embedding knowledge graphs and logical rules. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (2016)
6. Ji, G., He, S., Xu, L., Liu, K., Zhao, J.: Knowledge graph embedding via dynamic mapping matrix. In: Meeting of the Association for Computational Linguistics & the International Joint Conference on Natural Language Processing (2015)
7. Kejriwal, M., Szekely, P.: MYDIG: personalized illicit domain-specific knowledge discovery with no programming. *Future Internet* **11**(3), 59 (2019)
8. Yehuda, K., Robert, B., Chris, V.: Matrix factorization techniques for recommender systems. *Computer* **42**, 30–37 (2009)
9. Lin, H., Liu, Y., Wang, W., Yue, Y., Lin, Z.: Learning entity and relation embeddings for knowledge resolution. In: *Procedia Computer Science international Conference on Computational Science, ICCS 2017, 12–14 June 2017*, pp. 345–354. Zurich, Switzerland (2017)
10. Liu, P., et al.: HKDP: a hybrid knowledge graph based pediatric disease prediction system. In: *International Conference on Smart Health* (2016)
11. Galárraga, L., Teflioudi, C., Hose, K., Suchanek, F.M.: Fast rule mining in ontological knowledge bases with amie +. *VLDB J.* **24**, 707–730 (2015)
12. Mousselly-Sergieh, H., Botschen, T., Gurevych, I., Roth, S.: A multimodal translation-based approach for knowledge graph representation learning. In: *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pp. 225–234 (2018)
13. Nickel, M., Tresp, V., Kriegel, H.P.: A three-way model for collective learning on multi-relational data, pp. 809–816 (2011)
14. Ren, P., et al.: HMDFF: a heterogeneous medical data fusion framework supporting multimodal query. In: Siuly, Si., Wang, H., Chen, L., Guo, Y., Xing, C. (eds.) *HIS 2021. LNCS*, vol. 13079, pp. 254–266. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-90885-0_23
15. Schlichtkrull, M., Kipf, T.N., Bloem, P., Berg, R.V., Welling, M.: Modeling Relational Data with Graph Convolutional Networks. *The Semantic Web* (2018)
16. Wang, Q., Wang, B., Guo, L.: Knowledge Base Completion Using Embeddings and Rules. *AAAI Press* (2015)
17. Wang, Z., Zhang, J., Feng, J., Chen, Z.: Knowledge graph embedding by translating on hyperplanes. *Proc. AAAI Conf. Artif. Intell.* **28**(1) (2014)
18. Weston, J., Bordes, A., Yakhnenko, O., Usunier, N.: Connecting language and knowledge bases with embedding models for relation extraction. *IEEE* (2013)
19. Xie, R., Liu, Z., Jia, J., Luan, H., Sun, M.: Representation learning of knowledge graphs with entity descriptions. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 30 (2016)
20. Xie, R., Liu, Z., Luan, H., Sun, M.: Image-embodied knowledge representation learning. In: *AAAI Press* (2016)

21. Zhang, Y., Sheng, M., Zhou, R., Wang, Y., Dong, J.: HKGB: an inclusive, extensible, intelligent, semi-auto-constructed knowledge graph framework for healthcare with clinicians' expertise incorporated. *Information Processing & Management*, p. 102324 (2020)
22. Zhao, C., Jiang, J., Xu, Z., Guan, Y.: A study of EMR-based medical knowledge network and its applications. *Comput. Meth. Programs Biomed.* **143**, 13–23 (2017)