

Linking Physicians to Medical Research Results via Knowledge Graph Embeddings and Twitter

Afshin Sadeghi^{1,2} and Jens Lehmann^{1,2}

¹ Smart Data Analytics Group (SDA), University of Bonn, Germany

² Fraunhofer Institute for Intelligent Analysis and Information Systems, Sankt Augustin & Dresden, Germany

{sadeghi, jens.lehmann}@cs.uni-bonn.de

{afshin.sadeghi, jens.lehmann}@iais.fraunhofer.de

Abstract. Informing professionals about the latest research results in their field is a particularly important task in the field of health care, since any development in this field directly improves the health status of the patients. Meanwhile, social media is an infrastructure that allows public instant sharing of information, thus it has recently become popular in medical applications. In this study, we apply Multiple Distance Knowledge Graph Embeddings (MDE) to link physicians and surgeons to the latest medical breakthroughs that are shared as the research results on Twitter. Our study shows that using this method physicians can be informed about the new findings in their field given that they have an account dedicated to their profession.

Keywords: Knowledge Graph Embeddings · Social Media · Social Good · Health care · Twitter · Machine Learning

1 Introduction

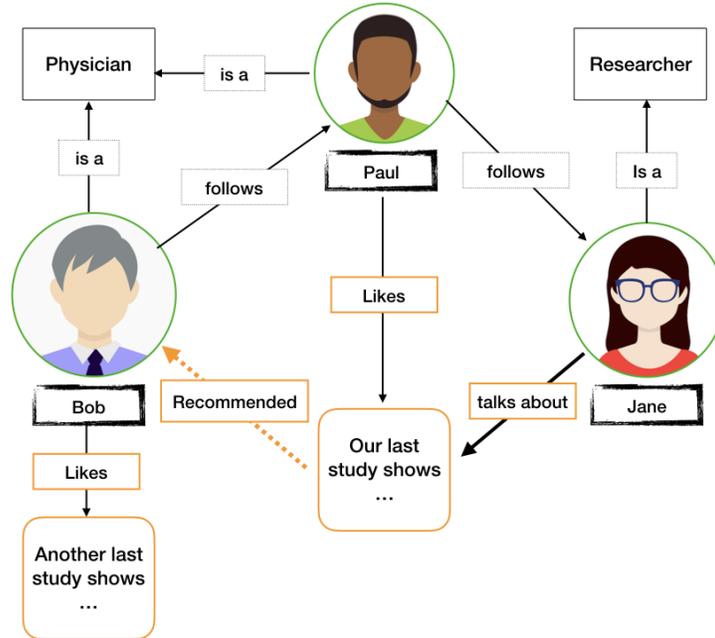
Twitter is a projection of the interactions of a society connected to the internet, which is in constant evolution. The dynamic aspect of this social media allows manifold applications. From the rise of social media, Twitter was used to measure campaign impacts, collect opinions, analyze trends and to study crisis. However, recently, its applications are more individualized. Particularly, because Twitter has become the most popular form of social media used for health care communication [15], and it is reshaping health care [10], it has become the center of many studies in the field of health care. For example, a study suggests Twitter for knowledge exchange in academic medicine [7] and it was argued that disease-specific hashtags and the creation of Twitter medical communities [13] has improved the uniformity of medical discussions. Another study is dedicated to the influence of specific medical hashtags on social media platforms [5].

Problem Statement: Pershad et al. [14] point out the potential of Twitter to reshape public health efforts, including disseminating health updates, sharing information about diseases. Especially, they emphasize on the role of Twitter

to make research advances more accessible for physicians. They argue that connecting researchers and clinicians is crucial and useful since clinicians can use new information they discover from this closer contact with researchers to guide decision-making about patient treatments in such a field that is in constant progress.

In this study, we target this problem by providing a method that suggests physicians and clinicians the recent research breakthroughs in their specialized field based on their current social activity. As the first step to reach this goal, we extract a subset of Twitter network and we generate a knowledge graph(KG) from the extracted data. Figure 1 depicts a schema of the KG with example user instances and the relations between them. In this figure, it is shown that our method recommends a Tweet of Jane, who is a researcher about her latest findings to Bob who is a surgeon. The method calculates a probability that such a Tweet will be useful to Bob based on his previous favored Tweets and the relation to other physicians that work in the same field.

Fig. 1. A schema of the medical professional knowledge graph on Twitter with example user instances. An orange dot line depicts a new link suggested by the proposed method.



We then apply an embedding method to predict links that are likely to be serving to the physicians. The proposed application is different from the user recommendation service of Twitter [9] which recommends users to follow or the

works that discover similar users [8]. Here, we focus our study and evaluation on suggesting related Tweets.

The main contribution of this study is the application of knowledge graph embeddings for a new field, the extraction of a KG from Twitter for the proof of concept and experiments targeted at physicians to help them stay up-to-date in their field.

2 Background

In a social network, a node of the graph represents a person while edges that link the nodes correspond to relationships between people. The edges are also called "connections", "links". Examples of social networks are graphs that describe Facebook and Twitter. Link prediction, in general, is the task of predicting whether a link exists between a given pair of nodes or not.

Definition 1: a knowledge graph is defined by (E, P, T) , a set of entities $e \in E$, a set of predicates $p \in P$ and a set of triples $t \in T$. A triple (e_i, p_k, e_j) is made of two entities and a predicate that connects them.

In a KG, two entities can be connected by several predicates. When describing a social network by a KG, nodes are translated to entities and links are translated to predicates, however, in KG an ontology usually specifies the class of describes what types of entities and predicates can construct a triple. A relational learning model usually learns the relations of a KG. Particularly, embedding models are a class of relational learning models that produce vector representations of the entities and predicates and predict the missing links.

Definition 2: link prediction in a KG means to predict the existence of a triple, i.e., whether a relation exists from two entities e_i and e_j and a k -th predicate.

3 Related Work

Classic link prediction methods on social media use graph properties of the social network or NLP feature of nodes to predict links between entities. For example, [11] is base solely on graph features and [3] uses a similar technique for the social networks in healthcare. Meanwhile, [1] uses common words to cluster and rank nodes and based on that predicts the closely-ranked nodes to be connected. Another Study [2] uses a combination of graph features and keyword matches to train classifiers(SVM, Naive Bayes, etc) to predict if a link exists between two nodes.

Most of the studies on link prediction of social networks focus on the problem of link existence. Where some methods attempt to find link weights and the number of links between the nodes [17]. An advantage of link prediction using KG

embedding is that the type of links are also predicted since these KG embedding models distinguish the type of links.

TransE [4] is an embedding model that is popular because of its simplicity and efficiency. It represents the entities in a KG by a relation between the vectors representing them. The score function describing these vectors in TransE is:

$$Score_{TransE} = \| h_i + r_i - t_i \|_p \quad (1)$$

where n refers to L_1 or L_2 norm and h_i and t_i are the vector representations of an entity and r_i is the vector representations of a predicate. For training, TransE uses margin ranking loss as the loss function. The following Section describes embedding a KG extracted from Twitter using MDE model.

4 KG Embeddings for Twitter Link Prediction

Knowledge graph embedding models usually generate a prediction based on their score function. Nickel et al. [12] suggests performing link prediction by comparing the score of a triple with some given threshold θ or by ranking the entries according to their likelihood that the link in question exists. We similarly use the Multiple-Distance Embedding (MDE) model [16]. In comparison to TransE, this model can learn several relational patterns and thus more can more accurately learn the hidden relation between the entities. Specifically, MDE can learn relations with symmetry, antisymmetry, transitive, inversion and composition patterns. The score function of this model is as follows:

$$Score_{MDE} = w_1 \| h_i + r_i - t_i \|_p + w_2 \| h_j + t_j - r_j \|_p + w_3 \| t_k + r_k - h_k \|_p - \psi \quad (2)$$

where $\psi \in \mathbb{R}^+$ is a positive constant. The loss function of this model is:

$$loss = \beta_1 \sum_{\tau \in \mathbb{T}^+} [f(\tau) - \gamma_1]_+ + \beta_2 \sum_{\tau' \in \mathbb{T}^-} [\gamma_2 - f(\tau')]_+ \quad (3)$$

where γ_1, γ_2 are small positive values and $\delta_0, \delta'_0 = 0$. $\beta_1, \beta_2 > 0$ are constraints. The given loss minimizes the score of the positive samples. Therefore, the smaller the score of a triple, that relation is more probable. Base on this property of the loss function, we define a measure to estimate the existence of a predicate such that the more probable triples are given a higher score.

We designate the division of the maximum score of a triple in the training set to the score of a triple A as the probability of the existence of A:

$$P_a = \frac{\max(Score_{training-triples})}{Score_A} \quad (4)$$

This definition is based upon the assumption that after the training, the model accurately predicates the triples of the training set.

Fig. 2. A sample of the extracted Tweets about the recent medical studies. Each row shows the content of one of the extracted Tweets.

- A research has shed light on the causes of #SpinalStenosis with ties connecting to the body's #genetic...
- New #research is proving that #stemcells can act as a living bandage for damaged tissue. https...
- Active #research is underway to make #hipreplacement surgeries safer and more effective. Insights...
- How new research on #stemcells could help patients in avoiding #kneereplacements...
- A recent study uses stem cells in the treatment of knee injuries...
- Collaborating with Rare Disease organizations to increase medical research ...
- A recent study looks at videotaped SCAs (sudden cardiac arrests) in athletes ...
- Inspire posts forward medical research in adverse drug reactions, new study shows...
- Immunotherapy empowers the immune system to fight cancer. New research on microbiome...
- What I told researchers trying to find a cure for my rare disease...
- More and more research indicates that a cancer diagnosis can cause PTSD, not only in patients...
- Paradigm shifting research!! New research shows how a blood test to detecting circulating tumor DNA...
- Patient-focused Drug Development or PFDD is an effort to increase the patients' voice in clinical research...
- New research indicates remission of type2diabetes is possible w/diet T2D ...
- A recent study conducted in Germany assessed post traumatic stress disorders in women diagnosed with brea...
- Parkinson's Disease (PD) affects ~10 million people worldwide. Recent research is close to...

The equation compares only the triples of the same type (with the same predicate). Thus, in predicting the triples for linking physicians to medical Tweets, we consider only triples with *like_research.Tweet_id* predicates.

To perform link prediction on Twitter, we train MDE over an extracted KG. In the following, we explain the procedure to extract the dataset from which we later generate a KG using it.

Knowledge Graph Extraction: We extract a set of Tweets about the latest medical studies using Python scripting and the Tweepy library ³. We filter our search by medical keywords and time in order to only obtain medical research related Tweets which were created from the beginning of the year 2019. Figure 4 shows a sample of the extracted Tweets. To keep the privacy of users, we removed the user and Tweet identifiers from the figure.

With the same tools, we search for Twitter users who are physicians, surgeons, nurses, and researchers in the medical fields which have written about these topics or favored such Tweets. Our continues inquiry which took 8 hours, provided us 5996 Twitter users. Between these users, the job title of 69 instances was deductible (researcher in the medical field or physician) based on the medical job titles in their profile descriptions. This step reduced the users to 69 instances. We then extract the relations among these users and the relations among the users and the gathered Tweets. We also extract the users who follow or are followed by these users so we gain the neighbors of these users in the social network.

We then generate a multi-relational knowledge graph from scraped data by converting these relations to triples. To create these triples, we first define an ontology for the social KG. This ontology includes five types of relations. Table 1 lists these relation types. We also anticipate two classes for users in the ontology. Table 2 presents these classes. The created knowledge graph (TW52) includes

³ <https://www.tweepy.org/>

4439 entities which are comprised of 1021 users and 3418 Tweets. The final constructed KG includes 4791 triples. The anonymised dataset is openly available for research purposes in <https://git.io/fj6h8>.

Table 1. Relation types in the Social Ontology

Relation Id	Relation
0	is_talking_about
1	is_followed_by
2	is_following
3	job_title_type_is
4	likes_research_tweet_id

Table 2. Class of users in the Social Ontology

Class Id	User Entity Class
0	job_title_medical_researcher
1	job_title_physician

5 Experiments

We set up two experiments. We firstly evaluate how well the MDE method performs on the social media dataset against a baseline in the task of link prediction. We then analyze the suggestion results of the model in different situations.

5.1 Performance Evaluation

We set up an experiment to evaluate the link prediction performance of MDE against TransE as the baseline.

Evaluation Setup: We dedicate 80 percent of the knowledge graph extracted from Twitter as the training dataset and set the rest as the test dataset. We randomly choose triples by uniform random selection to separate them for the test set. We perform ranking the score of each test triple against its versions with replaced head, and once with a replaced tail. We then compute the hit at $N(\text{hit}@N)$, mean rank(MR) and mean reciprocal rank (MRR) of these rankings. We set the vector size of TransE to 20 and choose the vector size of 10 for MDE. We use L_2 normalization to normalize their score function and train them by 700 iterations. For MDE, we set the hyperparameters as follows: $\gamma_1 = \gamma_2 = 3$ and $\psi = 1.2$.

Results: Table 3 lists the evaluation results of TransE and MDE on the extracted knowledge graph. Due to the sparsity of the graph, TransE gains very low ranking scores while MDE produces superior results for all the MR and MRR and $\text{hit}@N$ tests. The results suggest the positive influence of relation patterns learning in MDE.

Model	MR	MRR	Hit@1	Hit@3	Hit@10
TransE	1327	0.021	0.005	0.019	0.048
MDE	1287	0.148	0.071	0.161	0.332

Table 3. Results on Twitter extracted dataset (TW52). Better results are in bold.

5.2 Link Prediction Analysis:

In this section rather than studying the performance of the model, we establish an experiment to analyze the suggestion results of the model to find out whether it creates sound suggestions in different situations. We apply the model to learn on the constructed KG and then we use it to suggest the possible interesting research results for the physicians suggested. We then study the suggested results.

Considering the physicians in the KG and Tweets which include research results, we calculate the probability that such a Tweet is favorable for physicians using Equation 4. In our experiment, the hit@1 of the training triples was 99.8 percent, therefore, assuming maximum probability for the training triples in the formula holds for the experiment.

The observation of relations and entities in the KG shows that it is structured with the small world network patterns [18]. Particularly, it includes hub users and Tweets which are connected to other nodes with a number of links that greatly exceed the average degree in the network.

We select a subset of physicians in the KG and classify them according to their relation to other users and Tweets into 4 groups of 5 users. We particularly inspected their relation to hub users, which we call them User type A. Users of type A are followed by a large number of users(at least 200), they are active users and have favored variant Tweets. We consider also Users of type B which follow a small number of users(25) who are also physicians or researchers. Table 4 lists these groups of users and the mean of their probability to like a Tweet C that includes a research study. We consider two Tweets similar if their representative vectors have a small angle. These Tweets are usually favored by the same group of people.

Table 4. Mean Probability of linking to a Tweet C for users with different communities and liked Tweets

User group	Mean Probability of C
Users U that follow A. A and U like a Tweet similar to C	0.205
Users U that follow A. A likes a Tweet similar to C	0.134
Users U that follow B. A and B like a Tweet similar to C	0.975
New users U that still follow nobody and like no Tweet	0.127

It is observable from Table 4 that in the proposed model, users that follow a diverse group of users and topics, are less likely to be interested in an inquired

Tweet than those with less diverse connections. This effect is even stronger than when the user has liked a similar Tweet before. This suggests that the model performs better if a Twitter account is dedicated only to social communications related to her profession.

Additionally, the new users that have not favored any Tweet are expected the least among the users to favor a Tweet.

6 Discussion of the specificity of the problem

The proposed experiment in the study has two major components. The first is data extraction and KG construction section which we specify it to the problem by data cleaning and filtering the extracted data and creating an ontology specific to the physicians and research related tweets. The result of this part is TW52 knowledge graph which is sparse in comparison to the conventional benchmark datasets of embedding models, i.e., WordNet18 and FB15K.

The second part of the study is the MDE embedding model. Although MDE is a general method for the link prediction problems, the evaluations showed that it is capable of embedding the sparse dataset much better than the state-of-the-art TransE model. Therefore we consider both components appropriate for the proposed problem.

7 Conclusion

We proposed the usage of multiple-distance knowledge graph embeddings (MDE) to suggest Tweets about medical breakthroughs to physicians. We extracted a KG of medical research Tweets and their relations to the users which medical researcher or physicians.

We evaluated MDE against TransE as the baseline in a link prediction test for the social network KG. Our experiment shows the superior ranking performance of MDE over the baseline. We defined a probability for link suggestion and provided an analytic study for it. We thereby conclude that the model can be suggested to serve in connecting the physicians and the up-to-date advances in the medical studies. Considering the time constraints of physicians on social media [6], automating such suggestions can help physicians to find news and trends relevant for medical research results more easily and in less time.

As future work, it would be interesting to extend this study on a large scale and provide it as a live service. In addition, future studies may investigate the social effect of such application to find its effect benefits for patients besides the physicians.

8 Acknowledgements

This study is partially supported by project MLwin (Maschinelles Lernen mit Wissensgraphen, grant no. 01IS18050F)⁴ and Cleopatra (grant no. 812997). The

⁴ <https://mlwin.de/>

authors gratefully acknowledge financial support from the Federal Ministry of Education and Research of Germany (BMBF) which is funding MLwin and European Union Marie Curie ITN that funds Cleopatra, as well as Fraunhofer IAIS.

References

1. Adafre, S.F., de Rijke, M.: Discovering missing links in wikipedia. In: Proceedings of the 3rd international workshop on Link discovery. pp. 90–97. ACM (2005)
2. Al Hasan, M., Chaoji, V., Salem, S., Zaki, M.: Link prediction using supervised learning. In: SDM06: workshop on link analysis, counter-terrorism and security (2006)
3. Almansoori, W., Gao, S., Jarada, T.N., Elsheikh, A.M., Murshed, A.N., Jida, J., Alhajj, R., Rokne, J.: Link prediction and classification in social networks and its application in healthcare and systems biology. *Network Modeling Analysis in Health Informatics and Bioinformatics* **1**(1-2), 27–36 (2012)
4. Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., Yakhnenko, O.: Translating embeddings for modeling multi-relational data. In: Advances in neural information processing systems. pp. 2787–2795 (2013)
5. Brady, R., Chapman, S., Atallah, S., Chand, M., Mayol, J., Lacy, A., Wexner, S.: # colorectalsurgery. *British Journal of Surgery* **104**(11), 1470–1476 (2017)
6. Campbell, L., Evans, Y., Pumper, M., Moreno, M.A.: Social media use by physicians: a qualitative study of the new frontier of medicine. *BMC medical informatics and decision making* **16**(1), 91 (2016)
7. Choo, E.K., Ranney, M.L., Chan, T.M., Trueger, N.S., Walsh, A.E., Tegtmeyer, K., McNamara, S.O., Choi, R.Y., Carroll, C.L.: Twitter as a tool for communication and knowledge exchange in academic medicine: A guide for skeptics and novices. *Medical Teacher* **37**(5), 411–416 (2015). <https://doi.org/10.3109/0142159X.2014.993371>
8. Goel, A., Sharma, A., Wang, D., Yin, Z.: Discovering similar users on twitter. In: 11th Workshop on Mining and Learning with Graphs (2013)
9. Gupta, P., Goel, A., Lin, J., Sharma, A., Wang, D., Zadeh, R.: Wtf: The who to follow service at twitter. In: Proceedings of the 22Nd International Conference on World Wide Web. pp. 505–514. WWW '13, ACM, New York, NY, USA (2013). <https://doi.org/10.1145/2488388.2488433>, <http://doi.acm.org/10.1145/2488388.2488433>
10. Hawn, C.: Take two aspirin and tweet me in the morning: how twitter, facebook, and other social media are reshaping health care. *Health affairs* **28**(2), 361–368 (2009)
11. Martincic-Ipsic, S., Mocibob, E., Perc, M.: Link prediction on twitter. In: PloS one (2017)
12. Nickel, M., Tresp, V., Kriegel, H.P.: A three-way model for collective learning on multi-relational data. In: ICML. vol. 11, pp. 809–816 (2011)
13. Pemmaraju, N., Thompson, M.A., Qazilbash, M.: Disease-specific hashtags and the creation of twitter medical communities in hematology and oncology. In: Seminars in hematology. vol. 54, pp. 189–192. Elsevier (2017)
14. Pershad, Y., Hangge, P., Albadawi, H., Oklu, R.: Social medicine: Twitter in healthcare. *Journal of clinical medicine* **7**(6), 121 (2018)

15. Raghupathi, W., Raghupathi, V.: Big data analytics in healthcare: promise and potential. *Health information science and systems* **2**(1), 3 (2014)
16. Sadeghi, A., Graux, D., Shariat Yazdi, H., Lehmann, J.: Mde: Multi distance embeddings for link prediction in knowledge graphs. arXiv preprint arXiv:1905.10702 (2019)
17. Vartak, S.: A survey on link prediction. State University of New York, Binghamton, NY-13902, USA (2008)
18. Watts, D.J., Strogatz, S.H.: Collective dynamics of ‘small-world’ networks. *nature* **393**(6684), 440 (1998)