Improving Access to Science for Social Good

Mehdi Ali[1,2], Sahar Vahdati[1,3],
Shruti Singh[4], Sourish Dasgupta[4], and
Jens Lehmann[1,2]

[1] Smart Data Analytics Group (SDA), University of Bonn, Germany
{mehdi.ali|vahdati|jens.lehmann}@cs.uni-bonn.de
http://sda.tech
[2] Fraunhofer IAIS, St. Augustin and Dresden, Germany
{mehdi.ali|jens.lehmann}@iais.fraunhofer.de
http://iais.fraunhofer.de
[3] University of Oxford {sahar.vahdati}@cs.ox.ac.uk
[4] Rygbee Inc., Delaware, USA
{shruti.singh|sourish.dasgupta}@raxter.io

Abstract One of the major goals of science is to make the world socially a good place to live. The old paradigm of scholarly communication through publishing has generated enormous amount of heterogeneous data and metadata. However, most scientific results are not easy to discover, in particular those results which benefit social good and are also targeted at non-scientific people. In this paper, we showcase a knowledge graph embedding (KGE) based recommendation system to be used by students involved in activities aiming at social good. The recommendation system has been trained on a scholarly knowledge graph, which we constructed. The obtained results highlight that the KGEs successfully encoded the structure of the KG, and therefore, our system could provide valuable recommendations.

Keywords: Machine Learning · Knowledge Graph Embeddings · Social Good.

1 Introduction

People with different backgrounds ranging from school students to high profile professionals around the world are engaged in several initiatives such as political movements, environmental protection and fund-raising with the goal to achieve individual, community and society well-being [19]. One example is the *Fridays For Future* movement which was initiated by the young student Greta Thunberg to demonstrate against Swedens climate policy [36]. One of her main demands has been that the actions of the government of Sweden should become sufficient in order to comply with the essence of the Paris Agreement. Her initiative has quickly gained a lot of attention and initiated demonstrations all over Europe, and later in different countries around the world. Greta Thunberg is an illustrative example of a young school student who recognized some of the scientific findings with regards to the climate change and understood its importance for the social good: "We want politicians to listen to the scientists"[5], "Why should I be studying for a future that soon may be no more, when no one is doing anything to save that future?". While scientists increasingly have been called to share research findings about climate change [43], many other topics that are relevant to social good do not have a comparable media presence. For this reason, the information needs of activists that are non-experts may remain unsatisfied with regards to these topics. However, information technology and the digitization of scientific artifacts haves increased the amount of available scientific resources and offer a great potential to fulfill the information needs of activists that are concerned about social good. An overwhelming amount of scientific artifacts such as publications and their metadata have been made available independent of any geographical or temporal constraints on the web [6, 19, 47]. However, for non-experts the effective access to these artifacts is limited. While there are already existing services such as Google Scholar (GS) to explore and retrieve scientific publications, they alone are not sufficient to effectively fulfill the information needs of non-expert activists. One of the main reasons is the discrepancy between their search behaviors and the functionality of these services: GS expects specific search queries in order to provide relevant content on the first result page whereas non-experts (for instance undergraduate students) tend to use simple keyword or phrase queries, do not refine their search queries (e.g. by analyzing metadata), and usually ignore retrieved results beyond the first result page [10, 18]. In addition, search engines and services such as Google Scholar are not developed with the specific goal of providing access to content related to social good. Therefore, there is a need for a domain-specific system that can be effectively used by non-experts to access scientific content related to social good.

An approach to structure related knowledge that can be used to perform concept-based retrieval instead of string matching are knowledge graphs (KGs) [6, 22] which represent information as a set of triples of the form $(h, r, t) \in KG$ where h and t represent entities and r their relation. Recently, knowledge graph embeddings (KGEs) that encode the entities and relations of a KG into vector spaces while maintaining structural characteristics of the KG became popular. These embeddings can be used for several downstream machine learning tasks including recommendation systems.

In this paper, we present a recommender system that suggests for an entity of interest (i.e., publication, author, domain and venue) a set of related entities which helps users to effectively find relevant content related to the topic of social good from the large amount of available information. Our contributions are: i.) a KG that contains information about publications, domains, authors and venues. We focused on publications that are related to real-world problems such as climate change, marine litter, right movement and cyber security, ii.) a baseline recommender system that exploits KGEs to provide recommendations. We trained four different KGE models i.e., TransE, TransR, TransD and Com-

[5] https://www.fridaysforfuture.org/greta-speeches#greta_fullspeech_feb21_2019

plEx, and selected TransE to provide recommendations that have been manually evaluated. While the general approach can be transferred to different domains, the proposed recommender system is domain-specific.

In the following, we give an overview of the related work (Section 2), explain the KGE models that are relevant in the context of this work (Section 3), describe the process of creating our KG (Section 4), present our recommendation system (Section 5), explain our performed experiments (Section 6), discuss the limitations of our system and point out future work (Section 7), and finally, we give a short summary of this work (Section 8).

2 Related Work

Through the development of specialized search engines, digital libraries, databases and social networks for the scholarly domain, the availability of scientific artifacts and their metadata has been facilitated. Google Scholar[6] is an online search engine that has been realized in 2004 and enables users to search for both the printed and digital version of articles. Aminer [7] provides a faceted browsers on top of its mining service for researchers. *ResearchGate* [8] is a social network for researchers in which they can present their scientific profiles, their publications and interact with each other research. *Mendeley* [9] is a desktop service with a web program produced by Elsevier for managing and sharing research papers. There are several efforts to provide enhanced services by representing metadata of scholarly artifacts in a structured form. A crowd-sourcing platform for metadata management of scholarly artifacts is introduced in [39], and the representation of metadata in a semantic format is proposed in [3]. In Chi *et al.* [6] a knowledge graph and a metadata management systems for smart education is presented. However, most of these services either lack a systematic recommendation service or provide specialized suggestions based on user profiles. To the best of our knowledge (apart from dedicated journals and university libraries [44]) the domain of social science lacks a comprehensive and specialized knowledge graph with analytical and recommendation services on top. In a recent work, an embedding based recommendation system for books has been proposed. However, the recommendations are limited to one entity type, i.e. books. In this study and the follow up work, we aim to provide a comprehensive and domain-specific system in order to assist users in finding relevant artifacts of different types. Through the use of machine learning approaches, the system proposes recommendations that are beyond simple keyword-matching based recommendations.

[6] `https://scholar.google.de/`
[7] `https://www.aminer.org/`
[8] `https://www.researchgate.net/`
[9] `https://www.mendeley.com/`

3 Knowledge Graph Embeddings

Knowledge graph embedding models can be roughly divided into *translational distance models* and *semantic matching models.* Translational distance models compute the plausibility of a triple based on distance function (e.g. based on the Euclidean distance) and semantic matching models determine the plausibility of a triple by comparing the similarity of the latent features of the entities and relations [41]. In the following, we describe KGE models that are relevant in the context of this work, however many others have been proposed.

TransE An established *translational distance model* is TransE [5] that models a relation r as the translation from head entity h to the tail entity t:

$$h + r \approx t \tag{1}$$

To measure the plausibility of a triple following scoring function is defined:

$$f_r(h,t) = -\|h + r - t\| \tag{2}$$

The closer the embedding of the tail is to the sum of the head and relation embeddings, the higher is the probability that the triple is correct. Besides its simplicity, TransE is computational efficient, and can therefore be applied to large scale KGs. However, TransE is limited in modeling 1-N, N-1 and N-M relations. For this reason, several extensions have been proposed.

TransH TransH [41] is an extension of TransE that addresses the limitations of TransE in modeling N-M relations. Each relation is represented by a additional hyperplane, and the translation from the head to the tail entity is performed in the relation specific hyperplane. First, the head and tail entities are projected into the relation specific hyperplane:

$$h_\perp = h - w_r^\top h w_r \tag{3}$$

$$t_\perp = t - w_r^\top t w_r \tag{4}$$

where w_r is the normal vector of the relation specific hyperplane. After projecting the head and tail entity, the plausibility of the triple *(h,r,t)* is computed:

$$f_r(h,t) = -\|h_\perp + d_r - t_\perp\|_2^2 \tag{5}$$

where d_r is the relation specific translation vector lying in the relation specific hyperplane.

TransR TransR [41] is an extension of TransH that encodes entities and relations, in contrast to TransE and TransH, in different vector spaces. Similarly to TransH, each relation is represented by a matrix M_r that is used to project the entities into the relational specific space:

$$h_r = hM_r \tag{6}$$

$$t_r = tM_r \tag{7}$$

Consequently, the scoring function is defined as:

$$f_r(h,t) = -\|h_r + r - t_r\|_2^2 \tag{8}$$

TransD TransD [14] is an extension of TransR that uses fewer parameters than TransR and does not involve matrix-vector multiplications. Entities and relations are represented by two vectors, of which h, r, t encode the meanings of the entities/relations, and h_p, r_p, t_p are used to construct projection matrices that are used to project the entities in relation specific spaces:

$$M_{rh} = r_p h_p^T + I^{m \times n} \tag{9}$$

$$M_{th} = r_p t_p^T + I^{m \times n}, \tag{10}$$

where I is the identity matrix. These matrices are used to compute the projections of the head and tail entity:

$$h_\perp = M_{rh} h \tag{11}$$

$$t_\perp = M_{rt} t \tag{12}$$

Based on the projected entities, the score of the triple (h,r,t) is computed:

$$f_r(h,t) = -\|h_\perp + r - t_\perp\|_2^2 \tag{13}$$

RESCAL RESCAL [21] is a *semantic matching model* that represents each entity as a vector and each relation as a matrix, M_r. It uses the following scoring function:

$$f_r(h,t) = h^T M_r t \tag{14}$$

The relation matrix, M_r, encodes pairwise interactions between the features of the head and tail entities.

DistMult DistMult [45] simplifies RESCAL by allowing only diagonal matrices:

$$f_r(h,t) = h^T diag(r) t \tag{15}$$

where $r \in R^d$ and $M_r = diag(r)$. This is more computationally efficient compared to RESCAL.

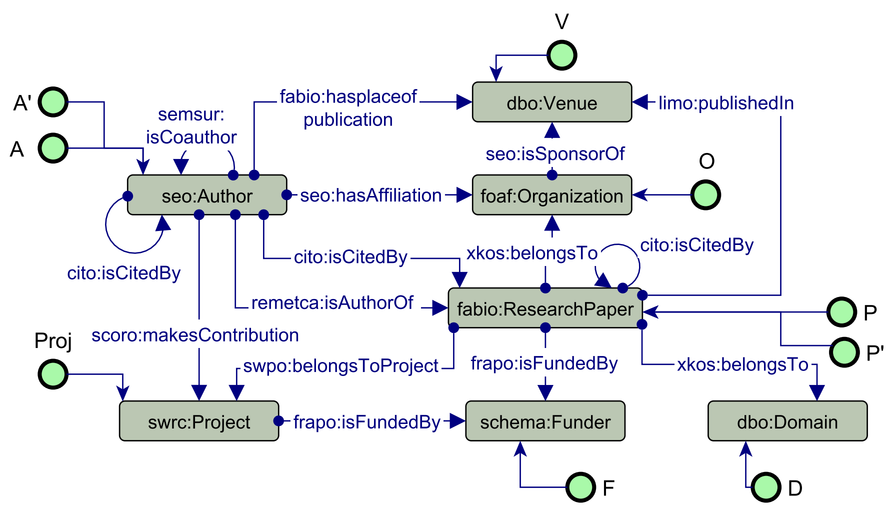

Figure 1: KG schema for social good.

ComplEx ComplEx [38] is an extension of DistMult in complex space. Considering the scoring function of DistMult (Equation 15) it can be observed that it has a limitation in representing anti-symmetric relations, because $h^T diag(r)t$ is equivalent to $t^T diag(r)h$. Equation 15 can be written in as the Hadamard product of h, r, t and t: $< h, r, t >= h * r * t$, where $h, r, t \in R^d$. The scoring function of ComplEx uses the Hadamard product in complex space, i.e. $h, r, t \in C^d$:

$$f_r(h,t) = \Re(\sum_{i=1}^{d} h_i * r_i * \overline{t_i}) \tag{16}$$

where $\Re(x)$ represents the real part of a complex number and $\overline{x}$ the conjugate of a complex number. It is straightforward to show that $f_r(h,t) \neq f_r(t,h)$, i.e. ComplEx is capable of modeling anti-symmetric relations.

4 Knowledge Graph Creation

As a first step, we created a scientific KG that gathers information relevant for social good which we used as a basis for providing recommendations. We defined the following requirements for the KG:

R1 The KG should contain publications with a focus on topics related to social good (e.g. climate change, social initiatives, political movements etc.)
R2 The KG should contain sufficient metadata to provide qualitative recommendations.

R3 The KG should be sufficiently large to allow conclusive insights about the applicability of modern machine learning methods.

To create the KG, the following steps were performed: i.) *Domain Conceptualization*, ii.) *Topic Conceptualization*, iii.) *Data collection* and *Data Curation*.

Domain Conceptualization Reusing already existing ontologies, we modeled a schema for KGs related to social good (see Figure 1) which will help researchers to get an overview of the domain, to define new KGs, and can be exploited by machines as an additional source of knowledge. Due to data availability, our KG currently, does not contain all types and relations described in the schema. Overall, seven core classes have been identified, namely *Papers*, *Venues*, *Authors*, *Organizations*, *Funders*, *Domains* and *Projects*. Furthermore, eight relationship types have been defined between these classes (see Figure 1).

Relation	Number of triples
authorOf	9090
isCoauthor	37326
hasPaperIn	6820
belongsToDomain	3998
isPublishedIn	3000
p_isCitedBy_p	355
a_isCitedBy_a	4388
a_isCitedBy_p	1225

Table 1: KG Statistics

Topic Conceptualization A list of topics have been collected from focal resources active in social good such as development program of United Nation[10] and sustainable development goals for 2030 [15] in addition to a systematic exploration on the Web. The topics have been short listed into four distinct categories as *climate change, political movements, marine/sea litter and cyber security*. This list have been used in the follow up steps of the

Data collection The data was collected using web crawlers of the RAx[11] platform which has reached to index metadata of 160+ million research paper. Based on the keywords, an exemplary dataset of 4004 matched papers has been extracted. The data was initially stored in JSON format (Listing 1.1), which we converted into a set of triples (Listing 1.2) representing our KG. KGs not only enable to represent data in form of triples, but also the metadata. For instance, for an entity representing a paper, we created triples of the form *(Paper1, belongsToDomain, Environmental Studies)* or *(Author1, authorOf, Paper1)*.

[10] `https://www.undp.org/`
[11] `https://raxter.io`

Listing 1.1: **Raw metadata 1.** JSON representation of original metadata.


```
*Paper1: {"title": T1,  "publication_venue": V1,
"citation_number": "50","doc_id": 7931d391-...dd32e50e8959,
"source_name": S1, "venue_id": 34544,
"raw_text": RW1, "authors": [A1, A2, A3],
"keywords":[ Social Science, ..., Climate Change],
"publisher": P1,  literature_type": Journal,
"source_url": [U1, U2, U3, U4],
"date": 2015-01-01,  "doi": 10.1016/...,
"references": [R1, R2, R3, R4 ]}*
*Paper2: { ...  }*
...
```


Listing 1.2: **TSV representation of metadata.** The KG in TSV format is used as input for the embedding models.


```
e9391a29-.... "belongsToDomain" Environmental studies.
Alan C. York "isCoauthor" G. J. Cary
Alan Manning "authorOf" 7604c5dc-...
Alark Saxena "hasPaperIn" Journal of Resources, Energy,
                        and Development
05607dab-... "p_isCitedBy_p" 635c28c3-...
05832950-... "isPublishedIn" International Womens Studies
A Allen "a_isCitedBy_a" Caroline S.E. Homer
```


5 System Description

The input to our workflow is a KG based on which a set of recommendations are computed in two major steps (Figure 2): i.) learning the KGEs and ii.) generating the recommendations based on the KGEs.

Learning the KGEs To learn the KGEs for our KG we make use of the software package PyKEEN [1] which integrates several KGE models. The learned embeddings encode structured knowledge represented in the KG. In the context of this work, we focused on the TransE, TransD, TransR and ComplEx model. The learned embeddings have been used as a basis for computing the recommendations.

Generating Recommendations For each seed entity that can be a selective entity in the KG (a publication, an author or a venue), the n nearest neighbors have been computed using the Euclidean norm (however, any similarity measure can be applied) and provided as recommendations. The recommendations for a seed publication can be for example the list of researchers who co-authored other publications with the authors in the seed publication, related publications or venues. The system is able to provide recommendations that represent n-hop dependencies in the KG. In turn, the provided recommendations can be used

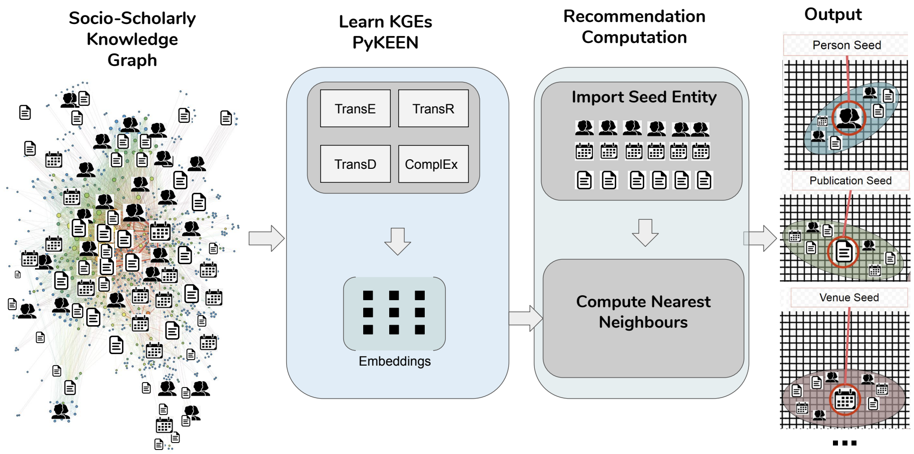

Figure 2: **A pipeline of recommendation services.**: i. embedded KG into latent feature space, ii.) filter publications based on KGEs, iii.) filter publications based on embeddings of their abstracts

as seed entities to access information that represents long-term dependencies in KG. Furthermore, the system is able to provide recommendations that are directly obvious due the capability of KGEs to capture global information of an entity and relation in a KG. The described steps don't require any complex traversing of the graph, instead, simple operations need to be applied on the learned embeddings.

6 Experiments

We evaluated four different KGE models (i.e., TransE, TransD, TransR and ComplEx) on the created KG. Afterward, we took one of the best performing models to provide the top n recommendations for a set of seed papers that have been manually evaluated. However, our approach be can applied on any type of seed entities.

6.1 Experimental Setup

We randomly split the initial KG into a training and test set where we took for each relation 10% of the triples which contain this relation as test triples. For each model, we performed a hyper-parameter optimization based on random search [11] and used *mean rank* and *hits@k* as evaluation metrics.

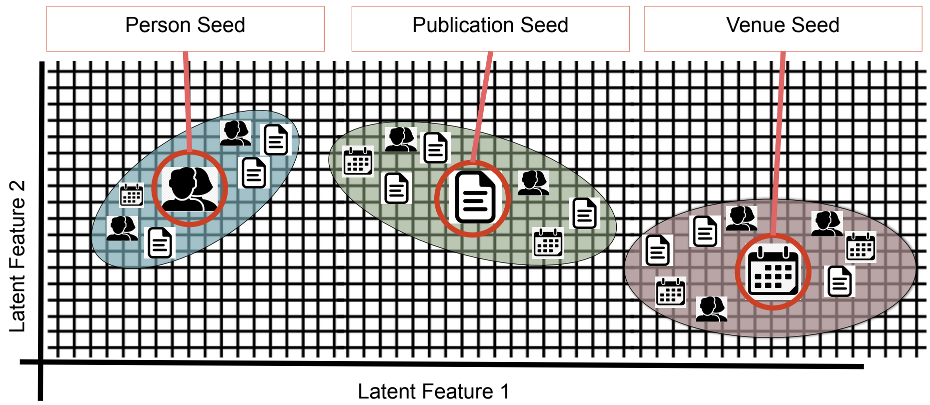

Figure 3: **Recommendations per seed entities.** For every seed entity type, a number of different recommendations are given.

6.2 Evaluation of the KGE Models

All the models have been trained based on the *open world assumption*, i.e., triples that are not part of the KG are not considered as non-existing, but as unknowns [22]. Therefore, we created artificial negative samples based on the negative sampling approach described by Bordes *et al.* [5]. For TransE, TransR and TransD the *margin ranking loss* that maximizes the distance between a positive and a corresponding negative triple [22] was applied, and for ComplEx and ComplEx* the *softplus loss* [38] was used. Furthermore, for all models except for ComplEx* one negative per each positive example, and for ComplEx* 10 negatives per each positive example were created in every forward step.

Model	**Mean Rank**	**Hits@10**
TransE [5]	90.63	91.22%
TransD [14]	443.84	6.43%
TransR [41]	397.02	43.63%
ComplEx [38]	267.91	83.00%
ComplEx* [38]	141.64	93.66%

Table 2: HPO results.

It can be observed that TransE, and ComplEx* performed very well (Figure 2). The high performance of TransE can be explained due to the fact that for most of the unique (subject, relation) and (relation, object)-pairs, there exists exactly one corresponding entity (subject/object) (77,72%/70.05% of the unique pairs). TransE even outperformed ComplEx when using only one negative for each positive example. However, for ComplEx only a few iterations of hyper-parameter optimization have been performed, and therefore, it is worthy to extend the hyper-parameter search. Similarly, the results of TransR and TransD might improve when applying a more extensive hyper-parameter search.

6.3 Recommendation of Related Information

Based on the results of the experiments, the TransE model has been selected to be used to compute and evaluate recommendations for a set of seed publications (Table 3 shows the titles, domains and venues of our seed papers). We choose TransE instead of ComplEx*, because it performs similarly and ComplEx(*) provides for each entity two vector representations those efficient combination should be investigated in more depth in a future work.

Domain	**Title**	**Venue**
Social Science	Cyber Bullying Detection Using Social and Textual Analysis [13]	System Analysis And Modeling
Social Science	Social Media, Indian Youth and Cyber Terrorism Awareness: A Comparative Analysis	Journal of Mass Communication and Journalism
Social Science	Expansion of Social Assistance: Does Politics Matter?	Economic and Political Weekly
Environmental studies	Reducing the impact of climate change	Bulletin of The World Health Organization
Environmental studies	General Chemistry Students' Understanding of Climate Change and the Chemistry Related to Climate Change	Journal of Chemical Education

Table 3: Selected seed publications.

Table 4 includes the validated recommendations for two of our seed papers from which one belongs to the domain of *Environmental Studies* and the second to the domain of *Social Science*. The recommendations are sorted according to their scores in descending order, i.e. the first recommendation received the highest score. For each recommended artifact, we performed a manual evaluation by looking them up in Google Scholar (the most used search engine for scholarly artifacts), and analysing their metadata. Among the recommendations there were obvious recommendations such as the authors of the seed papers (which we removed from the list of recommendations), irrelevant recommendations such as recommendation 11.) for the first seed paper, related publications (e.g. recommendation 2.) and 3.) for the first seed paper), co-authors of the authors of a seed paper (such as "Manfred Hauswirth") for the fist seed paper). Similar patterns can be detected in the recommendations for the second seed paper. The results highlight that relevant artifacts are recommended by the system. The recommendations indicate that the KGEs preserved the structure of the KG, for instance: i.) "Manfred Hauswirth" is a co-author of the authors of first seed paper, ii.) recommendation 1) that represents a publication, cites two of the authors of the seed paper ("Cory Andrew Henson", and "Vivek Kumar Singh").

Furthermore, it seems that the model has been able to distinguish entity types since the top recommendations usually represented publications for seed publications. While our evaluation approach indicates that the system is capable of providing relevant recommendations, involving external participants in the evaluation procedure will provide important insights regarding the effectiveness of our proposed system. In particular, we aim to perform a user study with expert and non-expert participants in order to analyse whether their information needs to topics related to social good can be fulfilled more effectively by using the proposed system. Because our work represents a preliminary work and such an evaluation requires an extensive preparation, we plan to target the described evaluation in a future work.

Recommendations for Cyber Bullying Detection Using Social and Textual Analysis [13]	**Type**
1) Physical-cyber-social computing: An early 21st century approach [33]	Paper
2) Physical cyber social computing for human experience [33]	Paper
3) Physical-Cyber-Social Computing (Dagstuhl Reports 13402) [34]	Paper
4) Transatlantic Social Politics: 1800-Present [30]	Paper
5) System-level design optimization for security-critical cyber-physical-social systems [48]	Paper
6) Cybermatics: Cyber–physical–social–thinking hyperspace based science and technology [24]	Paper
7) A cloud-edge computing framework for cyber-physical-social services [42]	Paper
8) Guest Editorial Data Mining in Cyber, Physical, and Social Computing [16]	Paper
9) Cyber-physical-social based security architecture for future internet of things [23]	Paper
10) Towards a politics of collective empowerment: Learning from hill women in rural Uttarakhand, India [31]	Paper
11) Manfred Hauswirth	Author
12) Payam M. Barnaghi	Author
13) Steffen Staab	Author
14) Markus Strohmaier	Author
15) Ramesh Jain	Author
16) Amit P. Sheth	Author
17) Social machine politics are here to stay [25]	Paper
18) IEEE Internet Computing Journal	Venue
Recommendations for: General Chemistry Students' Understanding of Climate Change and the Chemistry Related to Climate Change [40]	
1) Journal of Chemical Education	Venue
2) Marine Transportation and the Environment [17]	Paper
3) Stalinism and British Politics [37]	Paper
4) Piracy and the politics of social media [2]	Paper
5) Climate Change, Public Health and Sustainable Development: The Interlinkages [29]	Paper
6) Moisture dynamics in walls: response to micro-environment and climate change [12]	Paper
7) Diagenesis and Geochemistry of Sediments in Marine Environment [32]	Paper
8) Power, norms and institutional change in the European Union: The protection of the free movement of goods [8]	Paper
9) Adapting to climate change in Bangladesh: Good governance barriers [4]	Paper
10) Improving US Highway Safety: Have We Taken the Right Road? [26]	Paper
11) Climate change: the biggest challenge in the next decade?	Report
12) Social-Historical Transformations in Russia [20]	Paper
13) Fuller and Rouse on the legitimation of scientific knowledge [28]	Paper
14) High Politics, Low Politics, and Global Health [46]	Paper
15) Climate Change: A Serious Threat to Our Welfare and Environment [35]	Paper
16) Australian developments in marine science [7]	Paper
17) Pathways out of patronage politics: new roles for communities, new rules for politics in the Philippines [9]	Paper
18) Effects of climate change and variability on population dynamics in a long-lived shorebird [27]	Paper

Table 4: Recommendation for selected seed publications.

7 Limitations and Future Work

The approach presented in this paper represents a preliminary work that will be extended in future. Although the created KG contains already valuable information that we exploited to provide recommendations, it can benefit from several extensions. Currently, it contains only four entity and eight relationship types. We aim to augment this KG with additional information. In particular, we want to add entities that represent NGOs and other organizations, public speakers and events that are related to the topic of social good. Moreover, we want to provide major supporters/sponsors behind these organizations and events in order to provide more insights. Furthermore, we want to include relationship types that represent connections between (public/private) organizations to events and venues. The extended KG would contain more complex information that could be used to find non-obvious structures and to provide more diverse recommendations. For this work, we made use of KGE models that only consider triples of the form (h,r,t) where both h and t represent entities of the KG. However, there is a trend to develop multimodal KGE models that incorporate different types of information such as textual, numerical and visual information. In our future work, we plan to develop a multimodal KGE model in order to exploit textual information (e.g. abstracts of the papers) and numerical information (e.g. publication date, number of citations) which are available for our KG and might help to provide better recommendations.

Here, we provided recommendations by computing the nearest neighbors of a seed entity in the embedding space. Although this approach is easy to realize and provides interesting recommendations, it should serve as a baseline system for more sophisticated systems. As a next step, we aim to explore reinforcement learning based approaches in which feedback of the recommendations are taken into account during the training.

8 Conclusion

In this paper, we presented a socio-scholarly knowledge graph which contains information about scientific artifacts that are related to the topic of social good. A specific knowledge graph embedding-based recommendation system has been developed for this KG. The system provides recommendations for any given seed entity (publication, author, venue, domain) by returning related entities. Validated results show a great potential to leverage the system in broader scale of scholarly recommendations for active members of social good movements.

Acknowledgement This work is supported by the EPSRC grant EP/M025268/1, the WWTF grant VRG18-013, the EC Horizon 2020 grant LAMBDA (GA no. 809965), the CLEOPATRA project (GA no. 812997), and the German national funded BmBF project MLwin.

References

1. Ali, M., al., et. The KEEN Universe: An Ecosystem for Knowledge Graph Embeddings with a Focus on Reproducibility and Transferability. In: *International Semantic Web Conference*. 2019.
2. Almqvist, M. Piracy and the politics of social media. In: Social Sciences 5(3) (2016), p. 41.
3. Auer, S., Kovtun, V., Prinz, M., Kasprzik, A., Stocker, M., Vidal, M. E. Towards a Knowledge Graph for Science. In: *Proceedings of the 8th International Conference on Web Intelligence, Mining and Semantics*. ACM. 2018, p. 1.
4. Bhuiyan, S. Adapting to climate change in Bangladesh: Good governance barriers. In: South Asia Research 35(3) (2015), pp. 349–367.
5. 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*. 2013, pp. 2787–2795.
6. Chi, Y., Qin, Y., Song, R., Xu, H. Knowledge graph in smart education: A case study of entrepreneurship scientific publication management. In: Sustainability 10(4) (2018), p. 995.
7. Coffin, M. F. Australian developments in marine science. In: Advanced Earth and Space Science (2012).
8. Dimitrakopoulos, D. G. Power, norms and institutional change in the European Union: The protection of the free movement of goods. In: European Journal of Political Research 42(2) (2003), pp. 249–270.
9. Esguerra III, J., Villanueva, E. Pathways out of patronage politics: new roles for communities, new rules for politics in the Philippines. In: IDS Bulletin 40(6) (2009), pp. 13–21.
10. Georgas, H. Google vs. the library (part II): Student search patterns and behaviors when using Google and a federated search tool. In: portal: Libraries and the Academy 14(4) (2014), pp. 503–532.
11. Goodfellow, I., Bengio, Y., Courville, A. Deep learning. MIT press, 2016.
12. Hall, C., Hamilton, A., Hoff, W. D., Viles, H. A., Eklund, J. A. Moisture dynamics in walls: response to micro-environment and climate change. In: Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences 467(2125) (2010), pp. 194–211.
13. Huang, Q., Singh, V. K., Atrey, P. K. Cyber bullying detection using social and textual analysis. In: *Proceedings of the 3rd International Workshop on Socially-Aware Multimedia*. ACM. 2014, pp. 3–6.
14. Ji, G., He, S., Xu, L., Liu, K., Zhao, J. Knowledge graph embedding via dynamic mapping matrix. In: *Proceedings of ACL*. 2015, pp. 687–696.
15. Lee, B. X., Kjaerulf, F., Turner, S., Cohen, L., Donnelly, P. D., Muggah, R., Davis, R., Realini, A., Kieselbach, B., MacGregor, L. S. et al. Transforming our world: implementing the 2030 agenda through sustainable development goal indicators. In: Journal of public health policy 37(1) (2016), pp. 13–31.

16. Liu, J., Yan, Z., Vasilakos, A. V., Yang, L. T. Guest Editorial Data Mining in Cyber, Physical, and Social Computing. In: IEEE Systems Journal 11(1) (2017), pp. 194–196.
17. Meyer, M. D. Marine transportation and the environment. In: (2010).
18. Monchaux, S., Amadieu, F., Chevalier, A., Mariné, C. Query strategies during information searching: Effects of prior domain knowledge and complexity of the information problems to be solved. In: Information Processing & Management 51(5) (2015), pp. 557–569.
19. Mor Barak, M. E. The practice and science of social good: Emerging paths to positive social impact. In: Research on Social Work Practice (2018), p. 1049731517745600.
20. Nagapetova, A., Novikova, O., Pokhilko, A., Shmatko, A., Vetrov, Y. Social-Historical Transformations in Russia. In: International Journal of Humanities and Cultural Studies (IJHCS) ISSN 2356-5926 (2016), pp. 1439–1444.
21. Nickel, M., Tresp, V., Kriegel, H.-P. A three-way model for collective learning on multi-relational data. In: *Proceedings of the 28th international conference on machine learning (ICML-11)*. 2011, pp. 809–816.
22. Nickel, M., Murphy, K., Tresp, V., Gabrilovich, E. A review of relational machine learning for knowledge graphs. In: Proceedings of the IEEE 104(1) (2015), pp. 11–33.
23. Ning, H., Liu, H. Cyber-physical-social based security architecture for future internet of things. In: Advances in Internet of Things 2(01) (2012), p. 1.
24. Ning, H., Liu, H., Ma, J., Yang, L. T., Huang, R. Cybermatics: Cyber–physical–social–thinking hyperspace based science and technology. In: Future generation computer systems 56 (2016), pp. 504–522.
25. O'Hara, K. Social machine politics are here to stay. In: IEEE Internet Computing 17(2) (2013), pp. 87–90.
26. O'Neill, B. Improving US Highway Safety: Have We Taken the Right Road? In: TR News(239) (2005).
27. Pol, M. Van de, Vindenes, Y., Sæther, B.-E., Engen, S., Ens, B. J., Oosterbeek, K., Tinbergen, J. M. Effects of climate change and variability on population dynamics in a long-lived shorebird. In: Ecology 91(4) (2010), pp. 1192–1204.
28. Remedios, F. Fuller and Rouse on the legitimation of scientific knowledge. In: Philosophy of the Social Sciences 33(4) (2003), pp. 444–463.
29. Ruhil, R. Climate Change, Public Health and Sustainable Development: The Interlinkages. In: Indian Journal of Public Health Research & Development 7(3) (2016).
30. Scroop, D., Heath, A. Transatlantic Social Politics: 1800-Present. Palgrave Macmillan, 2014.
31. Sharma, D., Sudarshan, R. M. Towards a politics of collective empowerment: Learning from hill women in rural Uttarakhand, India. In: IDS Bulletin 41(5) (2010), pp. 43–51.
32. Sharma, G. D. Diagenesis and Geochemistry of Sediments in Marine Environment. In: AAPG Bulletin 50(3) (1966), pp. 634–634.

33. Sheth, A., Anantharam, P., Henson, C. Physical-cyber-social computing: An early 21st century approach. In: IEEE Intelligent Systems(1) (2013), pp. 78–82.
34. Sheth, A. P., Barnaghi, P., Strohmaier, M., Jain, R., Staab, S. Physical-Cyber-Social Computing (Dagstuhl Reports 13402). In: *Dagstuhl Reports*. Vol. 3. 9. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. 2014.
35. Singh, J., Agarwal, A., Singh, A. Climate Change: A Serious Threat to Our Welfare and Environment. In: ().
36. Thomson, D. et al. European left launches climate emergency manifesto. In: Green Left Weekly(1219) (2019), p. 16.
37. Thorpe, A. Stalinism and British politics. In: History 83(272) (1998), pp. 608–627.
38. Trouillon, T., Welbl, J., Riedel, S., Gaussier, É., Bouchard, G. Complex embeddings for simple link prediction. In: *International Conference on Machine Learning*. 2016, pp. 2071–2080.
39. Vahdati, S., Arndt, N., Auer, S., Lange, C. OpenResearch: collaborative management of scholarly communication metadata. In: *European Knowledge Acquisition Workshop*. Springer. 2016, pp. 778–793.
40. Versprille, A. N., Towns, M. H. General chemistry students understanding of climate change and the chemistry related to climate change. In: Journal of Chemical Education 92(4) (2014), pp. 603–609.
41. Wang, Q., Mao, Z., Wang, B., Guo, L. Knowledge graph embedding: A survey of approaches and applications. In: IEEE Transactions on Knowledge and Data Engineering 29(12) (2017), pp. 2724–2743.
42. Wang, X., Yang, L. T., Xie, X., Jin, J., Deen, M. J. A cloud-edge computing framework for cyber-physical-social services. In: IEEE Communications Magazine 55(11) (2017), pp. 80–85.
43. Wihbey, J., Ward, B. Communicating about climate change with journalists and media producers. In: *Oxford Research Encyclopedia of Climate Science*. 2016.
44. Wikipedia contributors. List of academic databases and search engines. [Online; accessed 28-June-2019]. 2019. `https://en.wikipedia.org/wiki/List_of_academic_databases_and_search_engines`.
45. Yang, B., Yih, W.-t., He, X., Gao, J., Deng, L. Embedding entities and relations for learning and inference in knowledge bases. In: arXiv preprint arXiv:1412.6575 (2014).
46. Youde, J. High Politics, Low Politics, and Global Health. In: Journal of Global Security Studies 1(2) (2016), pp. 157–170.
47. Zegura, E., DiSalvo, C., Meng, A. Care and the Practice of Data Science for Social Good. In: *Proceedings of the 1st ACM SIGCAS Conference on Computing and Sustainable Societies*. ACM. 2018, p. 34.
48. Zeng, J., Yang, L. T., Lin, M., Shao, Z., Zhu, D. System-level design optimization for security-critical cyber-physical-social systems. In: ACM Transactions on Embedded Computing Systems (TECS) 16(2) (2017), p. 39.