



HAL
open science

Knowledge Graph Refinement based on Triplet BERT-Networks

Armita Nassiri, Nathalie Pernelle, Fatiha Saïs, Gianluca Quercini

► **To cite this version:**

Armita Nassiri, Nathalie Pernelle, Fatiha Saïs, Gianluca Quercini. Knowledge Graph Refinement based on Triplet BERT-Networks. 2022. hal-03689800

HAL Id: hal-03689800

<https://hal.science/hal-03689800v1>

Submitted on 7 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Knowledge Graph Refinement based on Triplet BERT-Networks*

Armita Khajeh Nassiri¹[0000-0002-5734-0351], Nathalie
Pernelle^{1,2}[0000-0003-1487-393X], Fatiha Saïd¹[0000-0002-6995-2785], and Gianluca
Quercini¹[0000-0001-9195-1618]

¹ LISN, CNRS (UMR 9015), Paris Saclay University, France

² LIPN, CNRS (UMR 7030), University Sorbonne Paris Nord, France
firstname.lastname@lri.fr

Abstract. Knowledge graph embedding techniques are widely used for knowledge graph refinement tasks such as graph completion and triple classification. These techniques aim at embedding the entities and relations of a Knowledge Graph (KG) in a low dimensional continuous feature space. Unlike KG embedding methods and inspired by works built upon pre-trained language models, this paper adopts a transformer-based triplet network. It creates textual sequences from facts and fine-tunes a triplet network of pre-trained transformer-based language models, creating an embedding space that clusters the information about an entity or relation in the KG. It adheres to an evaluation paradigm that relies on efficient spatial semantic search technique. We show that this evaluation protocol is more adapted to a few-shot setting for the relation prediction task. Our proposed GilBERT method is evaluated on triplet classification and relation prediction tasks on multiple well-known benchmark knowledge graphs such as FB13, WN11, and FB15K. We show that GilBERT achieves better or comparable results to the state-of-the-art performance on these two refinement tasks.

1 Introduction

Large-scale knowledge graphs (KG) have gained popularity in the industry, like those built by Amazon, Google, and Microsoft, and in academia, such as YAGO and Wikidata. Knowledge graphs can be described as a structured repository of facts presented as triples (head, relation, tail). In other words, the facts are relations holding between two entities, for instance (Barack Obama, isMarriedTo, Michel Obama). More recently, KGs have become a natural resource for many AI downstream tasks such as question answering, entity disambiguation, and rule mining [4]. Although knowledge graphs contain millions of facts, they are considerably incomplete. Three different refinement tasks pertinent to this issue have attracted substantial attention in the literature: Triple classification, Relation Prediction, and Link Prediction. Most of the approaches proposed for tackling these tasks rely on embedding entities and relations of the KG into vectorized representations.

* Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Commonly, these embedding techniques [4] compute a defined scoring function that, given the fact, computes its plausibility. Hence, at the time of evaluation, depending on the nature of the task, the existence likelihood of unobserved facts (through substitution) should be calculated using the scoring function and then ranked. This causes additional overhead, especially for tasks such as link prediction. Others, like KG-BERT [19], can tackle the prediction task, e.g., relation prediction, as a classification task. Nevertheless, many relations in KGs are long-tail, meaning that very few facts with that relation exist in the KG. In such few-shot learning settings, the classification will be biased towards the data-rich relations, and the approach’s robustness can be hampered.

This paper proposes GILBERT, which employs a triplet architecture of a pre-trained language model such as BERT to cluster, in embedding space, all the facts regarding a relation or an entity. It overcomes the earlier shortcomings by fine-tuning BERT with a triplet loss that takes as input textual sequences created from facts. This way, depending on the refinement task, we force the information about an entity or a relation to be embedded close to each other, forming clusters. The evaluation protocol uses an efficient spatial semantic search technique, e.i., FAISS [5] to compute the closeness of the embedding of a test example to the point(s) in the embedding space. This reduces the overhead caused by substituting the candidates, and as shown in the experiments, it is more suited for a few-shot learning setting in the case of relation prediction. In addition, unlike translational and factorization-based models that need extra efforts to integrate any additional information such as entity descriptions, types, or paths to them [16] [14], it is trivial to do so in our proposed model. We also propose strategies that capture some of the semantics of the KG to generate harder negative information for the triplet network.

2 Related Work

A plethora of research is devoted to finding embeddings that model entities and relations of a KG in a low dimensional continuous vector space and to the proposition of scoring functions capable of capturing different kinds of relations. A recent comprehensive study overviews these techniques [4]. We can roughly categorize these models into three different categories. First, the translational-based models such as the TransE [1] and its more advanced extensions like TransH [15] and TransR [6], and more recently RotatE [12]. The translational-based models have proved to be successful from the beginning. TransE proposes one of the simplest and most prominent scoring functions that, given a fact, the head h and tail t embeddings are related by a translation vector r such that $h + r \approx t$. Secondly, the tensor factorization-based models with RESCAL [8] being one of the very earliest works adopting this technique. Other extensions DistMult [18], and ComplEx [13] instead of the tensor product, use dot product, and Hermitian dot product, respectively. Thirdly, the neural network-Based Models like ConvE[3] adopt CNNs with deeper scoring functions. Also, more recent transformer-based models such as KG-BERT[19], and R-MeN[7] have been proposed in the literature to solve different downstream tasks. To the best of our knowledge, KG-BERT is the only work that fine-tunes a pre-trained language model like BERT.

As pointed out in Section 1, unlike most of the previous works in the literature where the embeddings for entities and relations are learned, in our approach, we embed the training data into different clusters for entities or relations, depending on the downstream task. Following the indispensable contribution of transformers and pre-trained language models to NLP and the success of KG-BERT, it is natural to build upon BERT new methods that can be adapted to KGs. GilBERT adopts a triplet architecture that fine-tunes BERT (or any similar language model such as RoBERTa).

3 The proposed approach

Let \mathcal{G} be a knowledge graph containing a collection of facts represented as triples $\{(h, r, t)\} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$, where \mathcal{E} and \mathcal{R} denote the entity and relation sets, respectively. We propose GilBERT, an approach that fine-tunes BERT such that based on the refinement task, it clusters information about an entity (e.g., for triple classification task) or a relation (e.g., for relation prediction task). For some downstream tasks, not all triple parts (h, r, t) are known at the time of evaluation; e.g., the relation is unknown for relation prediction. Hence, to create the embedding space, we only use parts of the facts that are available at the time of evaluation so that the test data can be embedded the same way as during training. For the simplicity of notation, let a Partial fact PF be a textual sequence created out of a triple (h, r, t) by concatenating the entities or entities and relations. For instance, given the triple (Barack Obama, isMarriedTo, Michel Obama), PF_{rt} would be: "isMarriedTo Michel Obama", and PF_{ht} would be "Barack Obama Michel Obama". In all cases, and subject to availability, it is possible to replace the entity with its description to inject more background knowledge into the model.

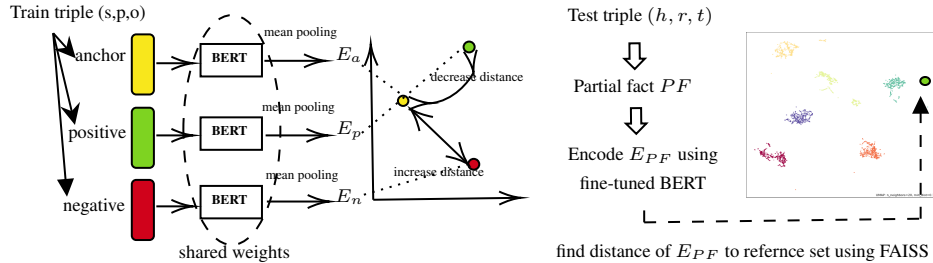


Fig. 1: (a) Overview of fine-tuning BERT using the partial facts of triples in \mathcal{G} . (b) Overview of evaluation, clusters illustrated using UMAP on part of FB15K dataset for the relation prediction task.

To cluster the training data, inspired by Sentence-BERT [9], we use a triplet network architecture of BERT that maps semantically similar facts near each other and dissimilar facts distant from one another. Thus it takes as input three partial facts PF , one as the anchor, one as positive (which is similar to the anchor), and one as negative (which is

dissimilar to the anchor). To obtain a fixed-sized embedding vector for the input PF , unlike KG-BERT, which uses the output of the [CLS] token produced in the last layer of BERT, a mean pooling layer is applied to BERT’s output layer. According to [9] and our experiments, the output of the pooling layer results in better embedding for the given input than using the embedding of the [CLS] token. This architecture clusters all the information about an entity (or relation). It is important to note that using PF s as inputs to the model allows us to use only the parts of a fact known at the time of evaluation (e.g., the head and the tail for a relation prediction task). This avoids the evaluation time overhead of previous models caused by creating all possible substitutions and ranking their plausibility based on the learned scoring function.

Figure 1 illustrates the fine-tuning of BERT through the triplet network that is comprised of three identical networks with shared parameters. The network is fed three input textual sequences: the partial facts of anchor, positive and negative, with the aim of anchor and positive ending up in the same cluster and anchor and negative in different ones. More specifically, the euclidean distance of the embedding of anchor E_a and positive E_p is reduced whilst that of the anchor E_a and negative E_n is increased. The triplet loss function is defined as:

$$\mathcal{L} = \max(0, \text{dist}(E_a, E_p) - \text{dist}(E_a, E_n) + \gamma),$$

where γ is the safety margin ensuring that the embedding of negative is at least γ further from the anchor than the positive.

To train the model, for each gold triple (h, r, t) in \mathcal{G} , n different input anchor, positive and negative PF s are created. It is important to create hard negative inputs so that the function \mathcal{L} is not trivial. In literature, to learn the scoring function that yields to embeddings, the negative examples are usually the corruption of a gold triple by replacing either its h or t . More sophisticated techniques such as [2] utilize background knowledge such as the schema of KG to better choose negatives. We consider the more frequently observed case where a schema is not available and try to capture the semantics of the KG when generating our input data. Here-under, we provide the details of data creation depending on whether the clustering is done based on relations or entities.

Clustering based on Relation. The anchor is the partial fact PF_{ht} . Each positive sample, $PF_{h_i t_i}$ is created from a fact drawn at random from the set $\{(h_i, r, t_i) \mid (h_i, r, t_i) \in \mathcal{G}\}$, and the negative samples $PF_{h_j t_j}$ from the set $\{(h_j, r', t_j) \mid (h_j, r', t_j) \in \mathcal{G} \wedge r' \in \text{close}_r\}$. So for generating hard negatives, the PF is created from a triple whose relation is semantically close to the relation of the anchor and positive. Denote the set F_r of all triples in \mathcal{G} having the relation r . The set close_r is chosen based on two criteria: (1) of all relations in \mathcal{R} , the percentage of tail entities shared with F_r , and (2) the percentage of head entities shared with F_r . The obtained percentages have the intuition of domain and range and close_r is the set of all relations whose computed sharing percentage is higher than a threshold. For instance, the $\text{close}_{\text{BornIn}} = \{\text{worksIn}, \text{diedIn}, \text{locatedIn}\}$.

Clustering Based on Entity. The clusters for entities are created using the entities appearing as head, and they can easily be extended to tail by generating reverse relations. The anchor is the partial fact PF_{rt} or randomly PF_{hrt} . The positive example is the partial fact created from a triple with the same head as anchor $\{(h, r_i, t_i) \mid (h, r_i, t_i) \in \mathcal{G}\}$. To create negative samples that are more difficult, we adhere to the idea that entities

that have the same type are semantically more related to each other. Hence we choose an entity that is likely to have the same type as the head entity h . We create PF_{rt_j} or $PF_{h_jrt_j}$ from the set $\{(h_j, r, t_j) \mid (h_j, r, t_i) \in \mathcal{G} \wedge (h, r, t_j) \notin \mathcal{G}\}$. This way, rather than setting apart a city from a person, which is too easy for the loss function, we disjoint two people who also share a property.

An overview of the evaluation pipeline for the relation prediction task is depicted in Figure 1. First, the partial fact of a test triple (h, r, t) is constructed in the same way as it has been done with the gold training inputs. The fine-tuned model is used to obtain the vectorized embedding E_{test} . The similarity, through Euclidean distance, of E_{test} to the embeddings of training inputs is used for the evaluation. Let the reference set Δ and Δ_e be the set of embeddings of all training anchors and the set of embeddings of anchors corresponding to entity e , respectively. For the triple classification, the distance of E_{test} is calculated with the reference set of its head entity Δ_h . And for relation prediction, the K closest points to E_{test} in Δ are used for choosing the most plausible relation. In both settings, to find the closest points to E_{test} in the corresponding reference sets, we rely on FAISS [5], a library for very efficient similarity search that trades precision for speed. More details of evaluation based on each refinement task will be outlined in Section 4.

4 Experimental Results

Datasets. In this paper, we adopt three widely used benchmark datasets in the literature: WN11 [11], FB13 [11], and FB15K[1]. The datasets are subsets of the real-world KGs, Wordnet, and Freebase, respectively. We have used the same training/validation/test splits as [11], and [1]. The test sets of WN11 and FB13 contain both positive and negative triples, whereas the FB15K dataset only contains correct triples. The statistics of the datasets are listed in Table 1.

Settings. In our experiments, we have fine-tuned the pre-trained RoBERTa-Base model for 4 epochs. We have used a batch size of 64, Adam optimizer with learning rate $2e-5$, and margin γ of 5. It should be noted that we have not used any external information such as paths or entity descriptions while training the model, and doing so can potentially enhance the results. All experiments have been run on a single machine with a processor 2.6GHz, 12 cores, 64 GB of RAM, and an Nvidia TITAN V GPU. The code and datasets are available at <https://github.com/armitakhn/gilbert/>.

Dataset	#R	#E	#Facts (Train/Valid/Test)		
WN11	11	38,696	112,581	2,609	10,544
FB13	13	75,043	316,232	5,908	23,733
FB15K	1345	14,951	483,142	50,000	59,071

Table 1: Statistics of datasets

Triple Classification. This task assesses whether a given triplet (h, r, t) is correct or not, i.e., binary classification on a triplet. The evaluation is done on two benchmark

datasets of WN11 and FB13. For this task, we cluster based on entities appearing as head in \mathcal{G} as explained in Section 3. The partial facts PF_{hrt} are used randomly with a 30% chance. For the WN11 dataset, the reverse versions of relations (t, r^{-1}, h) , are added to the KG to possibly overcome the cold start problem caused by some of the test data’s heads not appearing as the head to any fact in training. We create five different positive and five negative samples for each anchor in all datasets.

For the evaluation of a triple (h, r, t) , the embedding of PF_{hrt} ’s distance is calculated to the reference set Δ_h . We have exploited different decision strategies, MEAN, MAX, and MIN, to aggregate the computed distances to the reference set. If the aggregated distance is less than a threshold σ , then the triple is predicted positive and is predicted negative otherwise. The threshold σ is tuned on the validation data. Moreover, in our experiments, the best results were obtained by the MIN strategy and hence reported. The results of triple classification are shown in Table 2. The results for all methods have been taken from the cited papers in the Table, except for TransE, and DistMult whose results are taken from the papers [6] and [19], respectively. Our approach, GilBERT, has the best performance on average and on the FB13 dataset and is second best after KG-BERT for the WN11 dataset.

Method	WN11	FB13	Avg.
TransE [1]	75.9	81.5	79.2
TransH [15]	78.8	83.3	87.7
TransR [6]	85.9	82.5	83.9
DistMult [18]	87.1	86.2	86.7
DOLORES + ConvKB [14]	87.5	89.3	88.4
KG-BERT [19]	93.5	90.4	91.9
R-MeN [7]	90.5	88.9	89.7
GilBERT	92.7	92.5	92.6

Table 2: Triples classification Accuracy results (in %) reported on the WN11, FB13 test sets.

Relation Prediction This task is to predict the relation that holds between two given entities, i.e., $(h, ?, t)$. For this task, the clusters are created based on relations. The network is optimized to give small distances for embeddings of partial facts that belong to the same relation and large distances for relationally unrelated ones. We create five positive and five negative samples per anchor in the experiments.

For the evaluation, the correct relation r is determined given its head and tail. For this purpose the PF_{ht} ’s embedding E_{test} is derived. The k closest points in the embedding space (using all training points in the reference set Δ) to E_{test} are computed. The true relations of the closest points are known from the training data. We propose two different strategies for assigning a relation to our test point of interest. The strategy MIN assigns the relation of the closest point to E_{test} . And the strategy K-MODE chooses the relation that has most frequently appeared among the k -closest points of E_{test} .

The results of the relation prediction task on the FB15K dataset are reported in Table 3. All results are taken from their respective paper, except TransE and TransR

Method	MR	Hits@1
TransE [1]	2.5	84.3
TransR [6]	2.1	91.6
ProjE (listwise) [10]	1.2	95.7
TKRL [17]	1.7	92.8
DKRL(CNN) [16]	2.5	89.0
DKRL (CNN) + TransE [16]	2.0	90.8
KG-BERT [19]	1.2	96.0
GilBERT	1.3	92.0

Table 3: Relation prediction MR and HITS@1 results on FB15K dataset.

whose results are taken from [17]. The evaluation metrics used are Hits@1 (proportion of correct relations in top 1 ranking) and Mean rank MR . Note that the reported results are after filtering. If a suggested triplet is already known to be true, i.e., if it already exists in the train, validation, or test data of the knowledge graph, it is not correct to rank it among the suggestions. The best results with GilBERT, tuned on the validation data, were obtained using the K-MODE strategy with K being set to 10. The Hits@1 results obtained by our approach are comparable to the state-of-the-art, without integrating any external information, with a very competitive Mean Rank.

Moreover, as the introduction states, GilBERT is more adapted to a few-shot learning setting than KG-BERT, which tackles relation prediction as a classification task. We have investigated this by comparing the Hits@1 results of the two approaches on long-tail relation, trying relations that have less than 10, 15, 20, 25, and 30 facts belonging to them in the training data and hence covering 37.8, 43.2, 48.4, 51.8, and 54.9 percent of the relations in the training data respectively. The Hits@1 results on these proportions of relations in the test data are plotted in Figure 2 for KG-BERT and GilBERT. This plot shows that GilBERT is more robust at relation prediction for relations that have been seen only a few times during training.

To elaborate more, we look at 664 relations (almost half of the relations) with less than 20 facts, comprising 8.4% of FB15K instances. Our Hits@1 result on the test data involving these properties is 64% (best results obtained with MIN strategy) compared to KG-BERT, which gives a Hits@1 of 53%.

5 Conclusion

This paper proposed GilBERT, a new approach that fine-tunes a triplet network of pre-trained transformer-based language models by passing textual sequences created from facts. We explained how our evaluation paradigm that relies on a very efficient spatial semantic search technique has less overhead than previous approaches and is more suited for few-shot learning settings. Finally, we reported our results on two refinement tasks. In future work, we aim to try our approach for other KG refinement tasks such as link prediction, evaluating more benchmark datasets, and observing the impact of adding external information to training.




Fig. 2: Comparison of KG-BERT and GilBERT hits@1 results on long-tail relations of the training data.

Acknowledgements: This work has been supported by the project PSPC AIDA: 2019-PSPC-09 funded by BPI-France.

References

1. 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*. vol. 26. Curran Associates, Inc. (2013), <https://proceedings.neurips.cc/paper/2013/file/1cecc7a77928ca8133fa24680a88d2f9-Paper.pdf>
2. d’Amato, C., Quatraro, N.F., Fanizzi, N.: Injecting background knowledge into embedding models for predictive tasks on knowledge graphs. In: *The Semantic Web - 18th International Conference, ESWC 2021, Virtual Event, June 6-10, 2021, Proceedings. Lecture Notes in Computer Science*, vol. 12731, pp. 441–457. Springer (2021). https://doi.org/10.1007/978-3-030-77385-4_26, https://doi.org/10.1007/978-3-030-77385-4_26
3. Dettmers, T., Pasquale, M., Pontus, S., Riedel, S.: Convolutional 2d knowledge graph embeddings. In: *Proceedings of the 32th AAAI Conference on Artificial Intelligence*. pp. 1811–1818 (February 2018), <https://arxiv.org/abs/1707.01476>
4. Ji, S., Pan, S., Cambria, E., Marttinen, P., Yu, P.: A survey on knowledge graphs: representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems* (Apr 2021). <https://doi.org/10.1109/TNNLS.2021.3070843>
5. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus. *arXiv preprint arXiv:1702.08734* (2017)
6. Lin, Y., Liu, Z., Luan, H., Sun, M., Rao, S., Liu, S.: Modeling relation paths for representation learning of knowledge bases. In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. pp. 705–714. Association for Computational Linguistics, Lisbon, Portugal (Sep 2015). <https://doi.org/10.18653/v1/D15-1082>, <https://www.aclweb.org/anthology/D15-1082>

7. Nguyen, D.Q., Nguyen, T.D., Phung, D.: A relational memory-based embedding model for triple classification and search personalization. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL). pp. 3429—3435 (2020)
8. Nickel, M., Tresp, V., Kriegel, H.P.: A three-way model for collective learning on multi-relational data. In: ICML. pp. 809–816 (2011), https://icml.cc/2011/papers/438_icmlpaper.pdf
9. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics (11 2019), <https://arxiv.org/abs/1908.10084>
10. Shi, B., Weninger, T.: Proje: Embedding projection for knowledge graph completion. In: Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence. p. 1236–1242. AAAI’17, AAAI Press (2017)
11. Socher, R., Chen, D., Manning, C.D., Ng, A.: Reasoning with neural tensor networks for knowledge base completion. In: Advances in neural information processing systems. pp. 926–934 (2013)
12. Sun, Z., Deng, Z.H., Nie, J.Y., Tang, J.: Rotate: Knowledge graph embedding by relational rotation in complex space. In: International Conference on Learning Representations (2019), <https://openreview.net/forum?id=HkgEQnRqYQ>
13. Trouillon, T., Dance, C.R., Gaussier, É., Welbl, J., Riedel, S., Bouchard, G.: Knowledge graph completion via complex tensor factorization. *Journal of Machine Learning Research (JMLR)* pp. 1–38 (2017)
14. Wang, H., Kulkarni, V., Wang, W.Y.: {DOLORES}: Deep contextualized knowledge graph embeddings. In: Automated Knowledge Base Construction (2020). <https://doi.org/10.24432/C5SG64>, <https://openreview.net/forum?id=ajrveGQB10>
15. Wang, Z., Zhang, J., Feng, J., Chen, Z.: Knowledge graph embedding by translating on hyperplanes. p. 1112–1119. AAAI’14, AAAI Press (2014)
16. Xie, R., Liu, Z., Jia, J., Luan, H., Sun, M.: Representation learning of knowledge graphs with entity descriptions. *Proceedings of the AAAI Conference on Artificial Intelligence* **30**(1) (Mar 2016), <https://ojs.aaai.org/index.php/AAAI/article/view/10329>
17. Xie, R., Liu, Z., Sun, M.: Representation learning of knowledge graphs with hierarchical types. In: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence. p. 2965–2971. IJCAI’16, AAAI Press (2016)
18. Yang, B., Yih, S.W.t., He, X., Gao, J., Deng, L.: Embedding entities and relations for learning and inference in knowledge bases. In: Proceedings of the International Conference on Learning Representations (ICLR) 2015 (May 2015)
19. Yao, L., Mao, C., Luo, Y.: KG-BERT: BERT for knowledge graph completion. *CoRR abs/1909.03193* (2019), <http://arxiv.org/abs/1909.03193>