

Triple Classification for Scholarly Knowledge Graph Completion

Mohamad Yaser Jaradeh
L3S Research Center, Leibniz University Hannover
Hanover, Germany
jaradeh@l3s.de

Markus Stocker
TIB Leibniz Information Centre for Science and
Technology, Hanover, Germany
markus.stocker@tib.eu

Kuldeep Singh
Cerence GmbH
Zerotha Research, Aachen, Germany
kuldeep.singh1@cerence.com

Sören Auer
TIB Leibniz Information Centre for Science and
Technology, Hanover, Germany
auer@tib.eu

ABSTRACT

Scholarly Knowledge Graphs (KGs) provide a rich source of structured information representing knowledge encoded in scientific publications. With the sheer volume of published scientific literature comprising a plethora of inhomogeneous entities and relations to describe scientific concepts, these KGs are inherently incomplete. We present exBERT ($\epsilon\text{ksp3:rt}$), a method for leveraging pre-trained transformer language models to perform scholarly knowledge graph completion. We model triples of a knowledge graph as text and perform triple classification (i.e., belongs to KG or not). The evaluation shows that exBERT outperforms other baselines on three scholarly KG completion datasets in the tasks of triple classification, link prediction, and relation prediction. Furthermore, we present two scholarly datasets as resources for the research community, collected from public KGs and online resources.

CCS CONCEPTS

• **Information systems** → **Information extraction**; *Data mining*; • **Computing methodologies** → **Natural language processing**; *Information extraction*.

KEYWORDS

Scholarly Knowledge Graphs; Triple Classification; Link Prediction; Relation Prediction

ACM Reference Format:

Mohamad Yaser Jaradeh, Kuldeep Singh, Markus Stocker, and Sören Auer. 2021. Triple Classification for Scholarly Knowledge Graph Completion. In *Proceedings of the 11th Knowledge Capture Conference (K-CAP '21)*, December 2–3, 2021, Virtual Event, USA. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3460210.3493582>

1 INTRODUCTION

The proliferation of scientific literature creates new challenges in the research community, such as reproducibility crisis, duplication,

inefficiency, and a lack of transparency [42]. Thus, organizing scholarly knowledge is one of the most pressing tasks for solving current and upcoming societal challenges.

In general, Knowledge Graphs (KGs) have become central sources to structure knowledge and facts stored in so-called triples (*head, relation, tail*). These triples are used in downstream information extraction tasks such as entity linking, relation extraction, and question answering [12, 20]. Although large-scale KGs vary in terms of scope and coverage, they are often suffering from incompleteness and sparseness [22]. Therefore, KGs are required to be regularly updated and completed with missing information. The incompleteness of KGs motivates the knowledge graph completion task comprising several subtasks [10, 46]: i) *Link Prediction* aims to find missing head/tail entities in triples. ii) *Relation Prediction* predicts missing relations between two entities. KG completion methods calculate the plausibility of a triple via a scoring function to determine the validity of a knowledge graph triple [24]. These approaches can be broadly categorized into Knowledge Graph Embeddings (KGE) and Language Models (LM) [39]. KGE techniques such as TransE [9] and ConvE [14] learn entity and relation representation in a low dimensional vector space. However, due to their limitation of utilizing semantic information from text these methods produce different representations for the same entities in distinct triples [4]. Thus, LM techniques for KG completion emerged in an attempt to represent semantic information encoded in texts and produce contextualized embeddings [15].

Motivation. While KGs are a known solution for representing and managing for encyclopedic and common sense knowledge, e.g., in DBpedia [5] and Wikidata [38], the usage of KGs for scholarly knowledge is a rather new approach. Recently, researchers have focused on building scholarly KGs such as the Open Research Knowledge Graph [19] and MAG [16] to mitigate these previously mentioned challenges and to adhere to the FAIR data principles [43]. In contrast to the typical usage of KGs for encyclopedic and common sense knowledge, managing scholarly knowledge is significantly more challenging due to the heterogeneity, in-homogeneity and evolution of scholarly communication [18]. Scholarly KGs differ from encyclopedic purpose KGs because the encoded information is derived from the scholarly literature [13]. As a result, the scholarly KGs are ordinarily sparser than generic KGs because structured scholarly knowledge is more laborious to extract [19]. Furthermore, these KGs are ambiguous due to a lack of standard terminology used across the literature and poses domain-specific challenges for

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

K-CAP '21, December 2–3, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8457-5/21/12...\$15.00

<https://doi.org/10.1145/3460210.3493582>

KG completion task. Due to domain-specific challenges, the performance of existing KG completion methods such as KG-BERT and TransE are limited when applied to scholarly KG completion tasks (Section 4). The observed behavior is not surprising and is due to the peculiar entity and relation types of scholarly KGs. Hence, we suggest that existing KG completion approaches require additional task-specific context for the scholarly domain. We focus on the task of KG completion for the scholarly domain. Inspired by recent advancements in contextualizing language models [27], we argue that LMs can be utilized for scholarly KG completion tasks if fed with context derived from the scholarly literature. Our rationale for the choice is as follows: Language models such as SciBERT [6] are already trained in an unsupervised manner on large corpora of scholarly literature. Hence, adapting SciBERT for KG completion will allow us to inherit task-specific context. As such, we investigate

RQ1 *What is the impact of task-specific context on scholarly KG completion?*

Approach and Contribution. We model KG completion as a sequence classification problem by treating triples of the scholarly KG as sequences of natural language text. We propose exBERT, a contextualized approach for scholarly KG completion. In exBERT, we build on SciBERT [6] and adapt the underlying pre-trained language model on augmented sequences - augmented by appending types and synonyms of labels for triple elements - for predicting the plausibility of a triple, relation, or an entity. Our empirical results on the scholarly KG completion datasets provide superior performance on all three KG completion sub-tasks.

We provide the following three key contributions:

- i) exBERT a system that performs scholarly knowledge graph completion task, including the subtasks of triple classification, relation prediction, and link prediction.
- ii) Extensive evaluation on several datasets to show that proposed method achieves better results against various baselines for the respective tasks. Our proposed datasets, and code are publicly available¹.
- iii) We release two publicly available scholarly datasets¹ curated from scholarly KGs for the KG completion task.

2 RELATED WORK

The KG completion task utilizes different methods and techniques, which can be divided into the two main categories KGE and LM.

Knowledge Graph Embeddings (KGE): Knowledge graph embeddings are classified into several categories: translation, semantic matching, and neural network-based models [10]. Translational models use distance-based scoring functions to assess the plausibility of a triple (h, r, t) . Such models evaluate the plausibility by the distance between the head entity (h) and the tail entity (t) , which is typically done by performing a translation operation by the vector (r) [39]. Bordes et al. [9] proposed TransE as a representative model of translational graph embeddings; it uses a negative translational distance function for scoring. Work in [28] modified TransE for adapting it in scholarly domain. TransR [24] builds entities and relations embeddings by projecting them into different spaces, then building translations. Other models rely on

semantic matching, which internally uses similarity-based scoring functions. RESCAL [31], DistMult [45], and all their extensions are representatives of this category. Such models employ different scoring functions. For instance, DistMult uses a bilinear function $f(h, r, t) = \langle h, r, t \rangle$. These methods conduct KG completion using the structural information from the triples disregarding other external information, in particular entity types, logical rules, or textual descriptions. For example, Socher et al. [34] represented entities by averaging word embeddings extracted from their labels. Compared with these methods, our approach can learn context-aware text embeddings with pre-trained language models.

Language Models (LM): Language models that represent nodes and words can be split into two classes: feature-based and fine-tuning approaches. Widely used word embedding techniques such as Word2Vec [26] and its extensions aim at adopting features to learn context-independent word embeddings. GloVe [32] is another technique that strives to find a global representation for words based on specific features. FastText [8] is a technique that uses Skip-gram or Bag of Words to compute out-of-vocabulary word vector embeddings. More recently, Flair [1] embeddings try to generalize embeddings to be more context-aware. On the contrary, instead of relying on feature engineering, fine-tuning methods use the model architecture and parameters as a starting point for specific downstream NLP tasks. BERT [15] and Pegasus [48] are a few examples of pre-trained models that can be fine-tuned on a variety of NLP sub-tasks (e.g., Classification, or Summarization). Such models capture the deep semantic patterns in the natural language text and handle text ambiguity and variations. Pre-trained language models have also been used on knowledge graphs, where graph triples are converted into sentences using random walks and then used to train language models [17]. Other approaches aim to enhance BERT’s representations with entity semantics [49]. As such, these approaches concentrate on generating new entities and relations. Finally, KG-BERT [46] is a novel approach that utilizes a BERT transformer model to compute plausibility scores of triples based on names or descriptions of entities and relations.

3 APPROACH

We model the KG completion task as a sequence classification task to harness the richness and power of transformer language models. We rely on a BERT model to perform the classification via transforming the input knowledge graph triples into sequences of text with some extra tokens following the convention of BERT fine-tuning. Furthermore, we leverage a SciBERT model rather than a plain BERT as our core transformer model for the scholarly context.

3.1 Language Model for Scientific Text

SciBERT [6] is a state-of-the-art pre-trained language model that creates contextual representations using a multi-layer bidirectional transformer encoder architecture described by Devlin et al. [15]. SciBERT builds on the BERT [15] model and further trains it on scientific literature using 1.14 million scientific papers from Semantic Scholar [3]. The model also consists of 3.17 billion tokens (i.e., words). SciBERT is pre-trained on two tasks: masked language modeling and next sentence prediction. For next sentence prediction, the language model predicts whether two sequences of input

¹<https://github.com/YaserJaradeh/exBERT>

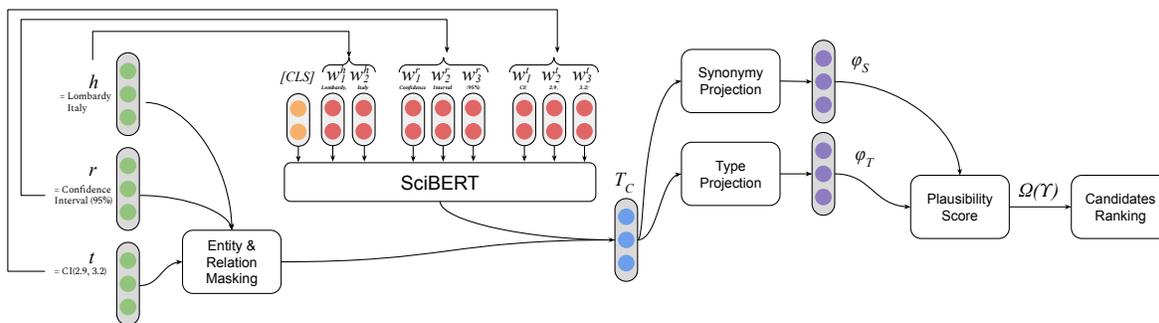


Figure 1: Leveraging triple classification to perform relation and link prediction tasks. T_C is the representation of the triple for the classification task. ϕ_S is the numerical representation after synonymy projection for relations. ϕ_T is the numerical representation after type projection for entities. $\Omega(\Upsilon)$ is the plausibility score of a triple .

(i.e., sentences) are consecutive in the text. For masked language prediction, SciBERT predicts suitably masked input tokens. Furthermore, regarding the fine-tuning phase, SciBERT is initialized with the weights and the parameters from the pre-training phase. Thus, these parameters are fine-tuned using labeled data from downstream natural language processing tasks (e.g., question answering, summarization, and token classification) [47].

3.2 Scholarly Model (exBERT)

In exBERT, we represent entities and relations of the knowledge graphs using their respective text labels. Furthermore, we augment these labels with type and synonymy projections. Types of entities are added to their labels and the synonyms of relations are included in the labels as well. These label sequences are given as an input sentence to our model for further fine-tuning. In order to estimate the plausibility score of triples, we arranged the sentences of (h, r, t) as individual sequences. In our case, a sequence represents a BERT compatible token sequence combined from two entities (*head* and *tail*) or a complete triple (*head, relation, tail*).

We extend the simple triples with extra information (i.e., types and label synonyms) to further leverage the KGs. We append to the triple the type label of each entity in the head or tail position, and augment the relation representation with label synonyms.

The workflow of exBERT is shown in Figure 1. exBERT performs a triple classification task to determine if a triple belongs to a KG. Furthermore, other tasks can be performed by leveraging the triple classification task, e.g., head/tail or relation prediction. The first token of every input sequence for exBERT is always a unique classification token [CLS]. Each entity in the triple and the relation are represented as a sentence containing a list of tokens v_1, \dots, v_α , $\alpha \geq 1$. For instance, *Lombardy Italy has a confidence interval (95%) of CI(2.9, 3.2)* has a head entity *Lombardy, Italy* which comprises two tokens $v_1^h = \text{Lombardy}$, $v_2^h = \text{Italy}$; the relation *Confidence Interval (95%)* comprises three tokens v_1^r, v_2^r, v_3^r ; and the tail *CI(2.9, 3.2)* comprises three tokens v_1^t, v_2^t, v_3^t . While constructing the sentences of entities and relations, a special token [SEP] is used to differentiate elements (i.e., components of the triple). The separation token indicates that the various elements of the sentence have different segment embeddings. However, the tokens of the head and tail entities share identical segment embeddings. Furthermore, when

various tokens occupy the same position – for the case of SciBERT, a position $1 \leq i \leq 512$ – they have the same position embeddings.

Token sequences are used as an input to the SciBERT model architecture, which is a multi-layer bidirectional transformer encoder based on the native BERT architecture [15]. The final hidden vector of the special [CLS] token is denoted as $C \in \mathbb{R}^H$, where H is the hidden state size in the pre-trained SciBERT model. Furthermore, the i -th token of the model’s input tokens is referred to as $v_i \in \mathbb{R}^H$. When fine-tuning the model to perform the triple classification task, a set of weights is created (i.e., classification layer weights) $W \in \mathbb{R}^{2H}$. A sigmoid function Ω is used to score a triple $\Upsilon = (h, r, t)$ and produce its class affiliation. Ω is defined as $\Omega^\Upsilon = \text{sigmoid}(CW^\Upsilon)$. Where $\Omega^\Upsilon \in \mathbb{R}^2$ is a two-dimensional real-valued vector with $\Omega_0^\Upsilon, \Omega_1^\Upsilon \in [0, 1]$ and $\Omega_0^\Upsilon + \Omega_1^\Upsilon = 1$.

We are then able to compute the cross-entropy loss using the scoring function Ω^Υ and predicted labels y^Υ as follows:

$$\Phi = - \sum_{\Upsilon \in \mathbb{K}^+ \cup \mathbb{K}^-} \left(y^\Upsilon \log(\Omega_0^\Upsilon) + (1 - y^\Upsilon) \log(\Omega_1^\Upsilon) \right) \quad (1)$$

where \mathbb{K}^+ is the positive triple set and \mathbb{K}^- is the negative triple set. y^Υ is the label of the triple (i.e., positive or negative label) $y^\Upsilon \in \{0, 1\}$. While the positive triple set contains the correct triples within the KG, the negative set \mathbb{K}^- is constructed by replacing the head entity h or tail entity t in a positive triple with a random entity \tilde{h} or \tilde{t} .

When corrupting the triple (i.e., creating the negative set), we make sure that the head or tail being replaced is not the correct one and ensure that the resulting corrupt triple does not belong to the positive set. Both relation and link prediction tasks use triple classification as an underlying task. The difference is in the way these tasks compose the input sentences (using only entities or the complete triple). Based on the findings of Yao et al. [46], we compose the input sequence with the two entities h and t to predict a relation r . KG-BERT empirical results suggest that predicting relations from the two entities using triple classification has higher performance. The other setting involving complete triples by curating negative samples with random relations \bar{r} does not yield a performance gain. Similarly to link prediction, the final hidden state vector C corresponds to the special classification token [CLS]. We highlight that relation prediction exBERT differs from link prediction exBERT (i.e., head and tail prediction) with the classification layer weights.

For relation prediction, the tasks fine-tune the weights $\bar{W} \in \mathbb{R}^{R \times H}$, whereby R is the number of all KG relations. The scoring function is $\bar{\Omega}^Y = \text{softmax}(C\bar{W}^v)$, whereby $\bar{\Omega}^Y \in \mathbb{R}^R$ is a \mathbf{R} -dimensional real vector with $\bar{\Omega}_i^Y \in [0, 1]$ and $\sum_i^R \bar{\Omega}_i^Y = 1$. Similarly, we compute the cross-entropy loss with the help of the scoring function $\bar{\Omega}^Y$ and the relation labels \bar{y}^Y as shown in Equation 2.

$$\Phi_r = - \sum_{Y \in \mathbb{K}^+} \sum_{i=1}^R \bar{y}_i^Y \log(\bar{\Omega}_i^Y), \quad \bar{y}_i^Y = \begin{cases} 1 & \text{if } r = i \\ 0 & \text{o/w} \end{cases} \quad (2)$$

where \bar{y}_i^Y is the relation class (i.e., indicator) for a positive triple Y and its value is conditional on the relation.

4 EVALUATION

We conduct our experiments and analysis in response to the overall research question **RQ1** (see section 1). As such, we also compare exBERT against approaches that do not come from a scholarly context. To understand the efficacy of exBERT for scholarly KG completion, we further divide our overall research question into three sub-questions: **RQ1.1** *What is the performance of exBERT for scholarly relation prediction?*, **RQ1.2** *What is the performance of exBERT for link prediction in scholarly KGs?*, and **RQ1.3** *What is the efficacy of exBERT in scholarly triple classification?*

Datasets. The scholarly KG domain is relatively new [19]. There is a scarcity of standard datasets to benchmark the performance of KG completion methods. Therefore, we created two datasets collected from available knowledge graphs and online resources.

i) **ORKG21:** A dataset of scholarly contributions extracted from the Open Research Knowledge Graph (ORKG) infrastructure [19]. The ORKG contains data on scientific contributions and publications curated by crowdsourcing that is complemented by automatic processes. We created the dataset based on a data dump provided by the ORKG in RDF [25] format. The resulting dataset contains many relationships because ORKG relies primarily on crowdsourcing for data entry and does not automate end-to-end extraction. As a consequence, multiple relationships are created by the different crowd members, but represent similar intentions, resulting in a significant number of relationships with slightly different representations.

ii) **PWC21:** A dataset from the online resource Papers-with-Code that describes papers in the field of machine learning, information extraction, and NLP along with their evaluation results. The PWC data is represented in tabular rather than graph structure. We transformed the raw data into RDF for broader use. The resulting dataset contains only a small set of relations because it focuses on certain aspects of research papers (i.e., evaluation results and metrics).

Since both datasets contain literals, we transformed them into entities by creating sequential IDs in the form of "/literal_num", similar to the Yago3-10 dataset [35]. In addition to the two datasets

Table 1: Statistics of datasets. We propose two new datasets (labeled with [*]) for the scholarly KG completion task.

Dataset	# Ent	# Rel	# Triples	# Train	# Dev	# Test
ORKG21*	226,210	2634	249,682	149,808	49,937	49,937
PWC21*	192,115	26	284,875	170,925	56,975	56,975
UMLS [7]	135	46	6,529	5,216	652	661

we created, we also used the public scientific dataset **UMLS** [7]. The ‘‘Unified Medical Language System’’ is an ontology for the medical field that describes technical medical concepts and their interlinked relations. Table 1 provides a summary of the employed datasets. If a dataset does not provide a class label (i.e., binary class) for training/testing triples, throughout our evaluation we considered every triple in the dataset as ground truth, and we corrupted triples to generate negative samples as explained in section 3.

Baselines. We compare exBERT with several state-of-the-art KGE methods to evaluate its efficacy with respect to the tasks:

- i) TransE [9] and its extensions TransH [41], TransD [21], TransR [24], TransG [44], and DistMult [45], which only rely on structural information of the knowledge graph to compute the embeddings.
- ii) NTN [34] and its simplified version ProjE [33].
- iii) ConvKB [29] which is a CNN model.
- iv) Other KGE models, specifically: AATE [4] and TEKE [40] that leverage textual information; Contextualized graph embeddings with KG2E [17]; Complex-values KG embeddings techniques RotatE [35] and ComplEX [36]; compositional vector-space embeddings HoLE [30]; Learning embeddings dependently with Simple [23]; v) BERT based approach KG-BERT by Yao et al. [46] which utilizes a BERT-Base model to perform KG completion.

It is noteworthy that the baselines vary by task because some KGE models do not perform tasks such as relation prediction.

Experimental Settings. SciBERT is the base transformer model used in exBERT. It has 12 layers, 12 self-attention heads, and $H = 768$ hidden layers. We used the Adam implementation for the optimizer. For fine-tuning the triple classification task, the batch size is 32, with a learning rate $5e - 5$ and a dropout rate of 0.1. For triple classification, we sample one negative triple for a positive triple to assure class balance for binary classification tasks. Furthermore, the number of epochs for triple classification is 3. We found no discernible improvement by increasing this number. For the link prediction task, we used 5 epochs, and we also sampled 5 negative samples for each positive triple (following the findings of Yao et al. [46]). Finally, for relation prediction, the number of epochs chosen was 20. For benchmarking exBERT and KG-BERT, we used a system running Ubuntu 20.04 with 128GB of RAM and 4x Nvidia A100 GPUs, each with 40GB vRAM. Other KGE baselines are trained on 8x Nvidia RTX 3090 GPUs, each with 24GB vRAM. KG embeddings are benchmarked using PyKeen [2].

Evaluation Metrics. Following the widely adapted metrics and inheriting evaluation settings from KG-BERT [46], we report the Mean Rank (MR) and the cut-off hit ratio (Hits@N) metrics on all the datasets for the link prediction ($N = 10$) and the relation prediction ($N = 1$) tasks. MR reports the average rank of all correct entities. Hits@N evaluates the ratio of correct entity predictions at a top N predicted results. Similar to KG-BERT, for the triple classification task we report accuracy. Classification accuracy is the number of correctly classified triples of the testing set divided by the total number of test triples. For the relation prediction task, we rank candidate relations by the scoring function $f(h, r, t)$ (cf. $\bar{\Omega}^Y$ in our approach formalization). Each correct test triple (h, r, t) is corrupted by replacing the relation with every other relation in the KG, with the exception of the relation itself, i.e., $r' \in R | r' \neq r$. Then

Table 2: Relation prediction results on all three datasets. Best results are indicated in bold font. The results listed in the table were all obtained by us. Techniques marked by [*] do not perform relation prediction by default. Hence, we employ triple classification to measure performance for relation prediction.

Method	ORKG21		PWC21		UMLS	
	MR	Hits@1	MR	Hits@1	MR	Hits@1
TransE [9]	1683	50.1	1232	49.5	166	23.3
TransH [41]	1544	51.8	1169	51.1	141	25.7
TransR [24]	1378	53.3	897	54.4	78	27.3
NTN* [34]	1189	58.4	754	59.2	22	29.9
KG2E* [17]	1301	57.9	773	58.3	21	30.0
ProjE [33]	522	74.6	212	78.5	6.42	41.8
PairRE [11]	206	82.1	59	88.6	4.25	63.5
KG-BERT [46]	15.37	92.8	1.51	96.7	1.21	87.2
exBERT	12.98	95.5	1.02	98.3	1.11	88.8

these candidates are ranked in descending order by their plausibility score we obtain from the triple classification task. Following Yao et al. [46] and Bordes et al. [9] we report the results under the filtered settings only, which means that all corrupted triples are removed from training, development, and testing sets before getting the ranking lists. In the link prediction task, we aim to predict the missing entity of the triple. Each correct test triple $T = (h, r, t)$ is corrupted by replacing either the head entity or the tail entity with every other entity in the KG $e \in E | e \neq h \wedge e \neq t$.

4.1 Experiment 1: Relation Prediction

Table 2 presents relation prediction results on all three datasets. Structural embeddings such as TransE, TransH, and TransR report limited performance due to the lack of contextualized information. KG2E and NTN perform slightly better than TransE and extensions with their density-based embeddings and neural model, respectively. PairRE outperforms the other embedding-based models with its paired encoding of relations and its capability of encoding various types of relations (e.g., symmetric, inverse). A common downfall of embedding-based techniques is that they learn identical embedding representations of entities and relations and do not account for the different meanings that words might have in various contexts. The limited performance in the scholarly domain clearly validates the observation. In contrary, KG-BERT and exBERT significantly outperform embedding-based models due to their own contextualized embeddings learned using a large corpus of unstructured text. However, task-specific context enhanced the ability of exBERT in predicting scholarly relations compared to KG-BERT, which successfully answers our first sub-question **RQ1.1**.

4.2 Experiment 2: Link Prediction

We report link prediction results in Table 3. Across all datasets, exBERT achieves significantly higher performance compared to all baselines. On the ORKG21 dataset, the majority of relations are symmetric relations. Translation-based methods report limited performance on ORKG21 due to their inability to infer symmetric connectivity patterns of a KG. Furthermore, the TransX family uses only the structure of the KG and does not induce context information or labels of entities, which results in limited performance. On

Table 3: Link prediction results on all three datasets. Best results are indicated in bold font. Results accompanied by asterisk* are reported by Yao et al. [46], other results were obtained by us and not from respective publications.

Method	ORKG21		PWC21		UMLS	
	MR	Hits@10	MR	Hits@10	MR	Hits@10
TransE [9]	2879	51.2	3176	60.8	1.84*	89.9*
TransH [41]	2811	52.5	2994	61.3	1.80*	99.5*
TransD [21]	2791	53.2	2135	68.8	1.71*	99.3*
ConvKB [29]	216	70.1	388	72.9	-	-
CompLEX [36]	713	65.3	456	72.7	2.59*	96.7*
HolE [30]	98	73.7	97	76.4	-	-
CompGCN [37]	2.84	84.6	4.02	82.7	-	-
Simple [23]	3.40	82.4	3.91	82.9	-	-
KG-BERT [46]	2.03	86.1	4.03	82.7	1.47*	99.0*
exBERT	1.80	87.4	2.11	84.2	1.97	98.9

PWC21, connectivity patterns are evenly distributed across all relation types (one-to-one, many-to-many, etc.), and the performance of exBERT does not drop. The majority of entities contain multiple relations on the UMLS dataset, and exBERT can successfully predict the missing links while reporting a slightly lower performance compared to best results. Embedding-based models continued to show limited performance for link prediction tasks due to non-standard characteristics of scholarly entities. From our empirical results, we conclude that the task-specific context fed into exBERT for the scholarly domain has positively impacted the performance across all KG completion tasks (successfully answering **RQ1.2**).

4.3 Experiment 3: Triple Classification

The objective of triple classification is to assign a score to each triple $T = (h, r, t)$ depending on whether or not it belongs to the underlying KG. Table 4 shows the classification results of exBERT against all other baselines on the benchmark datasets. exBERT outperforms all baseline KGE approaches confirming the effectiveness of our approach for the scholarly domain. We note that translation-based KGE such as TranE could not achieve high scores because

Table 4: Triple classification accuracy (in percent) for different embedding methods. The listed results were obtained by us and are not from the corresponding publications. Best results are indicated with bold font.

Method	ORKG21	PWC21	UMLS	Avg.
TransE [9]	77.6	77.3	78.1	77.7
TransH [41]	78.8	77.9	79.2	78.7
TransR [24]	81.4	81.3	81.9	81.5
TransD [21]	84.3	83.6	84.9	84.2
TransG [44]	85.1	84.7	85.2	85.0
TEKE [40]	84.8	84.2	84.9	84.6
KG2E [17]	79.6	78.8	79.7	79.4
DistMult [45]	86.2	86.2	86.8	86.4
ConvKB [29]	87.3	86.8	83.1	85.7
HolE [30]	87.3	87.3	88.2	87.6
NTN [34]	85.0	84.9	85.1	85.0
Simple [23]	89.7	89.2	89.1	89.3
KG-BERT [46]	95.1	93.3	89.7	92.7
exBERT	97.1	96.0	90.3	94.5

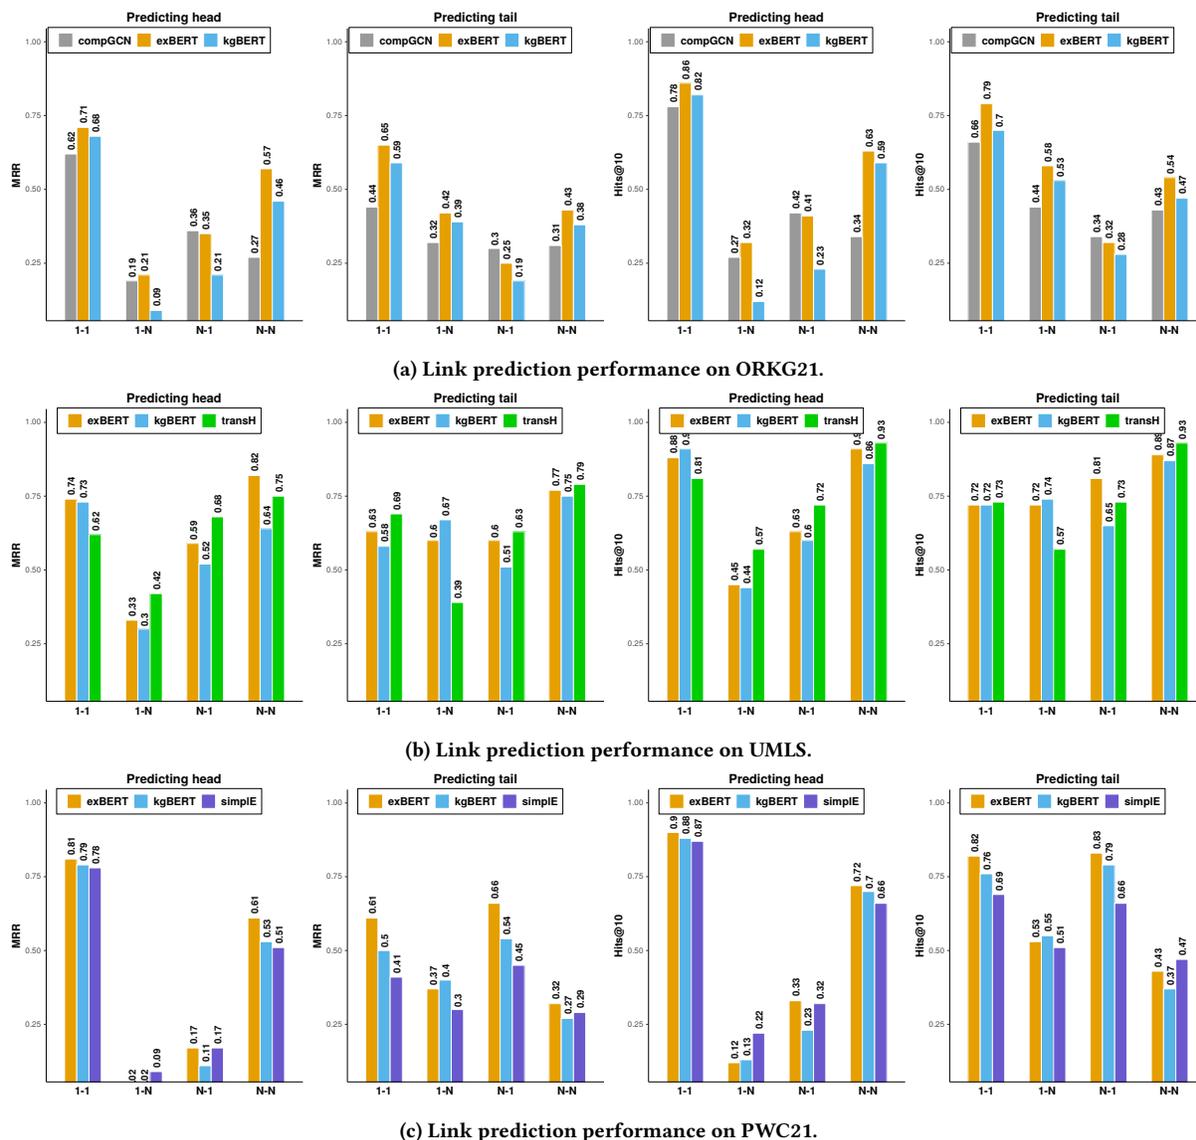


Figure 2: Link prediction performance for the top three performing baselines per dataset against exBERT for each relation type. We show two metrics MRR and Hits@10 for each relation type (1-1, 1-N, N-1, N-N). We observe that depending on the dataset and sub-task (head/tail prediction), performance varies per relation type. The fine-grained analysis provides a detailed overview of the strength and weakness of exBERT against best performing models.

of the cardinality of the knowledge graph relations (i.e., the existence of one-to-many, many-to-one, and many-to-many relations). However, its extensions, TransH, TransR, TransD, and TransG outperform TransE by introducing relation-specific parameters. The DistMult model performs relatively better than the Translation family. The CNN models, e.g., ConvKB, perform well, suggesting that a convolution model can capture global correlations among entities and relations in the scholarly domain. HolE performs proportionately well and comparably to ConvKB. SimpleE achieves the best results compared to the other embedding-based approaches. Finally, exBERT outperforms KG-BERT leveraging the underlying

semantic and contextualized knowledge of a transformer language model trained on scientific data.

Based on our observations, we identify the following reasons for the superior performance of exBERT for the triple classification task: Firstly, the triple classification task is akin to the next sentence prediction, for which SciBERT is pre-trained with large text corpora. The fine-tuning weights are already positioned for the inference of correlation among triple components. And mostly, contextualized embeddings for the scholarly domain are explicit with our approach via the hidden token vectors. This has positively

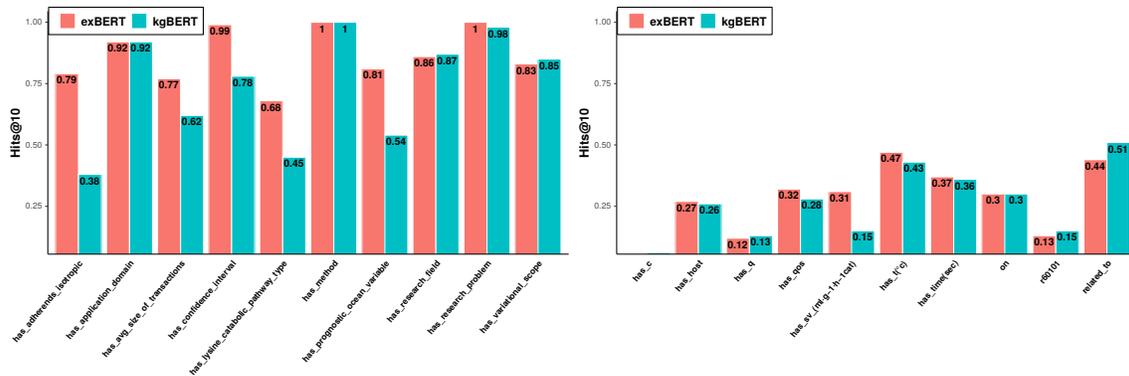


Figure 3: Illustration of the Hits@10 metric for the ten best (left plot) and worst (right plot) predicted relations in the ORKG21 dataset by exBERT and the second best baseline KG-BERT.

impacted the overall performance (successfully answering RQ1.3, and collectively answering RQ1).

4.4 Ablation Studies

Impact of relation type on the link prediction task. For each dataset, we show the metrics (Mean reciprocal rank (MRR) because it is in range [0,1] and Hits@10) for exBERT and the top-performing baselines for various relation types (i.e., one-to-one, one-to-many, many-to-one, many-to-many). Figure 2 presents the performance of this task for every relation type and all datasets. Based on the fine-grained analysis, we suggest that: i) On the ORKG21 dataset, exBERT maintains superior performance across all relation types for both metrics. CompGCN performs slightly better than KG-BERT and exBERT on many-to-one relation types. One possible reason is that CompGCN jointly embeds nodes and relations in a graph that permits the model to handle dense relations [37]. ii) On the UMLS dataset, exBERT suffers a performance drop in the head prediction that results in a lower performance observed in Table 3. iii) On the PWC21 dataset, it is interesting to observe that all three models suffer significant performance drops for predicting head entities in the one-to-many relation category. However, models maintain steady performance across all other relation categories.

Relation Prediction results for best/worst relations. To further comprehend how relations affect the overall performance of the relation prediction task, we select the best/worst ten performing relations in the ORKG21 dataset. We plot the evaluation metrics graphs for exBERT compared to the second best baseline KG-BERT. Figure 3 illustrates the individual performance. KG-BERT performs comparably to exBERT for the generic relation types such as *has_method*. However, for peculiar scholarly relations, such as *has_adreholds_isotrophic* and *has_prognostic_ocean_variable*, the performance of KG-BERT is limited due to missing task-specific context. Observed results validate our hypothesis to supplement exBERT with scholarly context for the KG completion task.

In Figure 3, most of the relations are quite ambiguous. Furthermore, in some cases, the scholarly context does not positively impact the performance of exBERT compared to KG-BERT. For instance, KG-BERT performs better than exBERT for the relation type

related_to. KG-BERT is trained on a large unstructured corpus from generic domains, and *related_to* is a commonly occurring relation between two real-world entities.

5 CONCLUSION AND FUTURE WORK

The hypothesis investigated in this paper was to study if task-specific context has an impact on scholarly KG completion task. For the same, we proposed exBERT and provided a set of experiments illustrating the positive impact of scholarly context encoded in exBERT for the KG completion task. We model the KG completion task as a sequence classification task, where we considered each KG triple as a set of sequences in a natural language. This allowed us to utilize SciBERT as an underlying model and adapt it for KG completion in the scholarly domain. We systematically studied the impact of our choices in the proposed approach. For instance, the ablation study demonstrates the effectiveness of scholarly context and provides insights on the strengths and weaknesses of exBERT. Albeit effective, exBERT is the first step of a larger research agenda. Based on our observations, we see the following open research questions in this domain: i) While scholarly context has positively impacted the performance, there are several relation types for which exBERT showed limited performance. One potential reason is the scarcity of training occurrences for certain relations such as *has_pos*. We believe that for unseen entities and relations, a zero-shot setting would be more suitable. Additional experiments are needed to verify the observation, and we point readers to this open research direction as the next step. ii) Considering that exBERT relies on labels and textual descriptions of an entity, data quality is a crucial aspect—as is typical for knowledge-intensive tasks. How the quality of contextual data impacts the performance of scholarly KG completion approaches is an open research direction. iii) Several recent approaches utilize KGs as an additional source of background knowledge to provide context in downstream tasks such as entity linking and relation extraction. How context derived from the scholarly KGs supplementing the textual context can help in effective KG completion is a further open question. iv) Our work focuses on English and developing multi-lingual KG completion approaches is a viable next step.

ACKNOWLEDGMENTS

This work was co-funded by the European Research Council for the project ScienceGRAPH (Grant agreement ID: 819536) and the TIB Leibniz Information Centre for Science and Technology. We thank Oliver Karras and Allard Oelen for their valuable feedback.

REFERENCES

- [1] Alan Akbik, Duncan Blythe, and Roland Vollgraf. 2018. Contextual String Embeddings for Sequence Labeling. In *COLING 2018, 27th International Conference on Computational Linguistics*. ACL, 1638–1649.
- [2] Mehdi Ali, Max Berrendorf, Charles Tapley Hoyt, Laurent Vermue, Sahand Sharifzadeh, Volker Tresp, and Jens Lehmann. 2021. PyKEEN 1.0: A Python Library for Training and Evaluating Knowledge Graph Embeddings. *Journal of Machine Learning Research* 22, 82 (2021), 1–6.
- [3] Waleed Ammar, Dirk Groeneveld, Chandra Bhagavatula, Iz Beltagy, Miles Crawford, Doug Downey, Vu Ha, Rodney Kinney, Lucy Wang, Chris Wilhelm, Zheng Yuan, Madeleine van Zuylen, and Oren Etzioni. 2018. Construction of the Literature Graph in Semantic Scholar. In *NAACL-HLT*. ACL, 84–91.
- [4] Bo An, Bo Chen, Xianpei Han, and Le Sun. 2018. Accurate Text-Enhanced Knowledge Graph Representation Learning. In *NAACL-HLT, Volume 1 (Long Papers)*. ACL, 745–755.
- [5] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. In *The semantic web*. Springer, 722–735.
- [6] Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. SciBERT: A Pretrained Language Model for Scientific Text. In *EMNLP-IJCNLP*. ACL, 3615–3620.
- [7] Olivier Bodenreider. 2004. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Research* 32, suppl_1 (01 2004), D267–D270.
- [8] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2016. Enriching Word Vectors with Subword Information. *arXiv preprint arXiv:1607.04606* (2016).
- [9] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating Embeddings for Modeling Multi-relational Data. In *Advances in Neural Information Processing Systems*, Vol. 26.
- [10] Hongyun Cai, Vincent W. Zheng, and Kevin Chen-Chuan Chang. 2018. A Comprehensive Survey of Graph Embedding: Problems, Techniques, and Applications. *IEEE* 30, 9 (Sept. 2018), 1616–1637.
- [11] Linlin Chao, Jianshan He, Taifeng Wang, and Wei Chu. 2021. PairRE: Knowledge Graph Embeddings via Paired Relation Vectors. *arXiv:2011.03798 [cs.CL]*
- [12] Wanyun Cui, Yanghua Xiao, Haixun Wang, Yangqiu Song, Seung-won Hwang, and Wei Wang. 2017. KBQA: Learning Question Answering over QA Corpora and Knowledge Bases. *Proc. VLDB Endow.* 10, 5 (Jan. 2017), 565–576.
- [13] Danilo Dessì, Francesco Osborne, Diego Reforgiato Recupero, Davide Buscaldi, Enrico Motta, and Harald Sack. 2020. AI-KG: An Automatically Generated Knowledge Graph of Artificial Intelligence. In *The Semantic Web – ISWC 2020*. Springer International Publishing, 127–143.
- [14] Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *NAACL-HLT*. ACL, 4171–4186.
- [16] Michael Färber. 2019. The microsoft academic knowledge graph: a linked data source with 8 billion triples of scholarly data. In *International Semantic Web Conference*. Springer, 113–129.
- [17] Shizhu He, Kang Liu, Guoliang Ji, and Jun Zhao. 2015. Learning to Represent Knowledge Graphs with Gaussian Embedding. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*. Association for Computing Machinery, 623–632.
- [18] Anthony JG Hey, Stewart Tansley, Kristin Michele Tolle, et al. 2009. *The fourth paradigm: data-intensive scientific discovery*. Microsoft research Redmond.
- [19] Mohamad Yaser Jaradeh, Allard Oelen, Kheir Eddine Farfar, Manuel Prinz, Jennifer D’Souza, Gábor Kismihók, Markus Stocker, and Sören Auer. 2019. Open Research Knowledge Graph. In *Proceedings of the 10th International Conference on Knowledge Capture*. ACM.
- [20] Mohamad Yaser Jaradeh, Markus Stocker, and Sören Auer. 2020. Question Answering on Scholarly Knowledge Graphs. In *International Conference on Theory and Practice of Digital Libraries*. Springer, 19–32.
- [21] Guoliang Ji, Shizhu He, Liheng Xu, Kang Liu, and Jun Zhao. 2015. Knowledge Graph Embedding via Dynamic Mapping Matrix. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th IJCNLP (Volume 1: Long Papers)*. ACL, 687–696.
- [22] Guoliang Ji, Kang Liu, Shizhu He, and Jun Zhao. 2016. Knowledge Graph Completion with Adaptive Sparse Transfer Matrix. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence (AAAI’16)*. 985–991.
- [23] Seyed Mehran Kazemi and David Poole. 2018. Simple Embedding for Link Prediction in Knowledge Graphs. In *Advances in Neural Information Processing Systems*, Vol. 31.
- [24] Yankai Liu, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion (AAAI’15).
- [25] Frank Manola, Eric Miller, Brian McBride, et al. 2004. RDF primer. *W3C recommendation* 10, 1-107 (2004), 6.
- [26] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and Their Compositionality. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2 (NIPS’13)*. 3111–3119.
- [27] Isaiah Onando Mulang, Kuldeep Singh, Chaitali Prabhu, Abhishek Nadgeri, Johannes Hoffart, and Jens Lehmann. 2020. Evaluating the impact of knowledge graph context on entity disambiguation models. In *CIKM*. 2157–2160.
- [28] Mojtaba Nayyeri, Sahar Vahdati, Jens Lehmann, and Hamed Shariat Yazdi. 2019. Soft marginal transe for scholarly knowledge graph completion. *arXiv preprint arXiv:1904.12211* (2019).
- [29] Dai Quoc Nguyen, Tu Dinh Nguyen, Dat Quoc Nguyen, and Dinh Phung. 2018. A Novel Embedding Model for Knowledge Base Completion Based on Convolutional Neural Network. In *NAACL-HLT, Volume 2 (Short Papers)*. ACL, 327–333.
- [30] Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016. Holographic embeddings of knowledge graphs. In *AAAI*, Vol. 30.
- [31] Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A Three-Way Model for Collective Learning on Multi-Relational Data (*ICML’11*). Omnipress, 809–816.
- [32] Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global Vectors for Word Representation. In *EMNLP*. ACL, 1532–1543.
- [33] Baoxu Shi and Tim Wengner. 2017. ProjE: Embedding Projection for Knowledge Graph Completion (AAAI’17). 1236–1242.
- [34] Richard Socher, Danqi Chen, Christopher D. Manning, and Andrew Y. Ng. 2013. Reasoning with Neural Tensor Networks for Knowledge Base Completion (*NIPS’13*). 926–934.
- [35] Zhiqing Sun, Zhi-Hong Deng, Jian-Yun Nie, and Jian Tang. 2018. RotatE: Knowledge Graph Embedding by Relational Rotation in Complex Space. In *International Conference on Learning Representations*.
- [36] Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. Complex embeddings for simple link prediction. In *International Conference on Machine Learning*. PMLR, 2071–2080.
- [37] Shikhar Vashishth, Soumya Sanyal, Vikram Nitin, and Partha Talukdar. 2020. Composition-based Multi-Relational Graph Convolutional Networks. In *International Conference on Learning Representations*.
- [38] Denny Vrandečić and Markus Krötzsch. 2014. Wikidata: A Free Collaborative Knowledgebase. *Commun. ACM* 57, 10 (Sept. 2014), 78–85.
- [39] Quan Wang, Zhendong Mao, Bin Wang, and Li Guo. 2017. Knowledge Graph Embedding: A Survey of Approaches and Applications. *IEEE Transactions on Knowledge and Data Engineering* 29, 12 (2017), 2724–2743.
- [40] Zhigang Wang, Juanzi Li, Zhiyuan Liu, and Jie Tang. 2016. Text-enhanced representation learning for knowledge graph. In *IJCAI*. 4–17.
- [41] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge Graph Embedding by Translating on Hyperplanes (AAAI’14). 1112–1119.
- [42] Karen E White, Carol Robbins, Beethika Khan, and Christina Freyman. 2017. Science and engineering publication output trends: 2014 shows rise of developing country output while developed countries dominate highly cited publications. *National Center for Science and Engineering Statistics InfoBrief* (2017), 1–7.
- [43] Mark D Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E Bourne, et al. 2016. The FAIR Guiding Principles for scientific data management and stewardship. *Scientific data* 3, 1 (2016), 1–9.
- [44] Han Xiao, Minlie Huang, and Xiaoyan Zhu. 2016. TransG: A Generative Model for Knowledge Graph Embedding. In *ACL (Volume 1: Long Papers)*. ACL, 2316–2325.
- [45] Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and Li Deng. 2014. Embedding entities and relations for learning and inference in knowledge bases. *arXiv preprint arXiv:1412.6575* (2014).
- [46] Liang Yao, Chengsheng Mao, and Yuan Luo. 2019. KG-BERT: BERT for Knowledge Graph Completion. *arXiv:1909.03193 [cs.CL]*
- [47] Haiou Zhang, Hanjun Zhao, Chunhua Liu, and Dong Yu. 2020. Task-to-Task Transfer Learning with Parameter-Efficient Adapter. In *Natural Language Processing and Chinese Computing*, Xiaodan Zhu, Min Zhang, Yu Hong, and Ruifang He (Eds.). Springer International Publishing, 391–402.
- [48] Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. PEGASUS: Pre-training with Extracted Gap-sentences for Abstractive Summarization. In *ICML*. PMLR, 11328–11339.
- [49] Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. ERNIE: Enhanced Language Representation with Informative Entities. *ACL*, 1441–1451.