

Automated Mining of Leaderboards for Empirical AI Research

Salomon Kabongo¹[0000–0002–0021–9729], Jennifer D’Souza²[0000–0002–6616–9509],
and Sören Auer^{1,2}[0000–0002–0698–2864]

¹ L3S Research Center, Leibniz University of Hannover, Hannover, Germany
kabenamualu@l3s.de

² TIB Leibniz Information Centre for Science and Technology Hannover, Germany
{jennifer.dsouza, soeren.auer}@tib.eu

Abstract. With the rapid growth of research publications, empowering scientists to keep oversight over the scientific progress is of paramount importance. In this regard, the LEADERBOARDS facet of information organization provides an overview on the state-of-the-art by aggregating empirical results from various studies addressing the same research challenge. Crowdsourcing efforts like PapersWithCode among others are devoted to the construction of LEADERBOARDS predominantly for various subdomains in Artificial Intelligence. Leaderboards provide machine-readable scholarly knowledge that has proven to be directly useful for scientists to keep track of research progress. The construction of LEADERBOARDS could be greatly expedited with automated text mining.

This study presents a comprehensive approach for generating LEADERBOARDS for knowledge-graph-based scholarly information organization. Specifically, we investigate the problem of automated LEADERBOARD construction using state-of-the-art transformer models, viz. Bert, SciBERT, and XLNet. Our analysis reveals an optimal approach that significantly outperforms existing baselines for the task with evaluation scores above 90% in F1. This, in turn, offers new state-of-the-art results for LEADERBOARD extraction. As a result, a vast share of empirical AI research can be organized in the next-generation digital libraries as knowledge graphs.

Keywords: Table mining · Information extraction · Scholarly text mining · Knowledge graphs · Neural machine learning.

1 Introduction

Our present rapidly amassing wealth of scholarly publications [29] poses a crucial dilemma for the research community. A trend that is only further bolstered in a number of academic disciplines with the sharing of PDF preprints ahead (or even instead) of peer-reviewed publications [13]. The problem is: *How to stay on-track with the past and the current rapid-evolving research progress?* In this era of the publications deluge [45], such a task is becoming increasingly infeasible even within one’s own narrow discipline. Thus, the need for novel

technological infrastructures to support intelligent scholarly knowledge access models is only made more imminent. A viable solution to the dilemma is if *research results* were made skimmable for the scientific community with the aid of advanced knowledge-based information access methods. This helps curtail the time-intensive and seemingly unnecessary cognitive labor that currently constitute the researcher’s task of searching just for the results information in full-text articles to track scholarly progress [9]. Thus, *strategic reading* of scholarly knowledge focused on core aspects of research powered by machine learning may soon become essential for all users [43]. Since the current discourse-based form of the *results* in a static PDF format do not support advanced computational processing, it would need to transition to truly digital formats (at least for some aspects of the research).

In this regard, an area gaining traction is in empirical Artificial Intelligence (AI) research. There, LEADERBOARDS are being crowdsourced from scholarly articles as alternative machine-readable versions of performances of the AI systems. LEADERBOARDS typically show quantitative evaluation results (reported in scholarly articles) for defined machine learning tasks evaluated on standardized datasets using comparable metrics. As such, they are a key element for describing the state-of-the-art in certain fields and tracking its progress. Thus, empirical results can be benchmarked in online digital libraries. Some well-known initiatives that exist to this end are: PAPERSWITHCODE.COM [4], NLP-Progress [3], AI-metrics [1], SQUaD explorer [6], Reddit SOTA [5], and the Open Research Knowledge Graph.³

Expecting scientists to alter their documentation habits to machine-readable versions rather than human-readable natural language is unrealistic, especially given that the benefits do not start to accrue until a critical mass of content is represented in this way. For this, the retrospective structuring from pre-existing PDF format of results is essential to build a credible knowledge base. Prospectively, machine learning can assist scientists to record their results in the LEADERBOARDS of next-generation digital libraries such as the Open Research Knowledge Graph (ORKG) [27]. In our age of the “deep learning tsunami,” [38] there are many studies that have used neural network models to improve the construction of automated scholarly knowledge mining systems [36,12,7,28]. With the recent introduction of language modeling techniques such as transformers [44], the opportunity to obtain boosted machine learning systems is further accentuated.

In this work, we empirically tackle the LEADERBOARD knowledge mining machine learning (ML) task via a detailed set of evaluations involving a large dataset and several ML models. The LEADERBOARD concept varies wrt. the domains or the captured data. Inspired by prior work [24,26,40], we define a LEADERBOARD comprising the following three scientific concepts: 1. *Task*, 2. *Dataset*, and 3. *Metric* (TDM). However, this base LEADERBOARD structure can be extended to include additional concepts such as method name, code links, etc. In this work, we restrict our evaluations to the core TDM triple. Thus, constructing

³ <https://www.orkg.org/orkg/benchmarks>

a LEADERBOARD in our evaluations entails the extraction of all related TDM statements from an article. E.g., (Language Modeling, Penn Treebank, Test perplexity) is a LEADERBOARD triple of an article about the ‘Language Modeling’ *Task* on the ‘Penn Treebank’ *Dataset* in terms of the ‘Test perplexity’ *Metric*. Consequently, the construction of comparisons and visualizations over such machine-interpretable data can enable summarizing the performance of empirical findings across systems.

While prior work [24,40] has already initiated the automated learning of LEADERBOARDS, these studies were mainly conducted under a single scenario, i.e. only one learning model was tested, and over a small dataset. For stakeholders in the Digital Library (DL) community interested in leveraging this model practically, natural questions may arise: *Has the optimal learning scenario been tested? Would it work in the real-world setting of large amounts of data?* Thus, we note that it should be made possible to recommend a technique for knowledge organization services from observations based on our prior comprehensive empirical evaluations [28]. Our ultimate goal with this study is to help the DL stakeholders to select the optimal tool to implement knowledge-based scientific information flows w.r.t. LEADERBOARDS. To this end, we evaluated three state-of-art transformer models, viz. Bert, SciBert, and XLNet, each with their own respective unique strengths. The automatic extraction of LEADERBOARDS presents a challenging task because of the variability of its location within written research, lending credence to the creation of a human-in-the-loop model. Thus, our LEADERBOARD mining system will be prospectively effective as intelligent helpers in the crowdsourcing scenarios of structuring research contributions [42] within knowledge-graph-based DL infrastructures. Our approach called ORKG-TDM is developed and integrated into the scholarly knowledge organization platform Open Research Knowledge Graph (ORKG) [27].

In summary, the contributions of our work are:

1. we construct a large empirical corpus containing over 4,500 scholarly articles and LEADERBOARD TDM triples for training and testing LEADERBOARD extraction approaches;
2. we evaluate three different transformer model variants in experiments over the corpus and integrate these into the ORKG-TDM LEADERBOARD extraction platform;
3. in a comprehensive empirical evaluation of ORKG-TDM we obtain scores of 93.0% micro and 92.8% macro F1 outperforming existing systems by over 20 points.

To the best of our knowledge, our ORKG-TDM models obtain state-of-the-art results for LEADERBOARD extraction defined as (*Task*, *Dataset*, *Metric*) triples extraction from empirical AI research articles. Thus ORKG-TDM can be readily leveraged within KG-based DLs and be used to comprehensively construct LEADERBOARDS with more concepts beyond the TDM triples. Our data⁴ and code is made publicly available.⁵

⁴ <http://doi.org/10.5281/zenodo.5105798>

⁵ <https://github.com/Kabongosalomon/task-dataset-metric-nli-extraction>

2 Related Work

Organizing scholarly knowledge extracted from scientific articles in a Knowledge Graph, has been viewed from various Information Extraction (IE) perspectives.

Digitalization based on Textual Content Mining. Building a scholarly knowledge graph with text mining involves two main tasks: 1. scientific term extraction and 2. extraction of scientific or semantic relations between the terms.

Addressing the first task, several dataset resources have been created with scientific term annotations and the term concept typing to foster the training of supervised machine learners. For instance, the ACL RD-TEC dataset [22] annotates computational terminology in Computational Linguistics (CL) scholarly articles and categorizes them simply as *technology* and *non-technology* terms. The ScienceIE SemEval 2017 shared task [10] annotates the full text in articles from Computer Science, Material Sciences, and Physics domains for *Process*, *Task* and *Material* types of keyphrases. SciERC [37] annotates articles from the machine learning domain with six concepts *Task*, *Method*, *Metric*, *Material*, *Other-ScientificTerm* and *Generic*. The STEM-ECR [18] corpus annotates *Process*, *Method*, *Material*, and *Data* concepts in article abstracts interdisciplinarily across ten STEM disciplines.

For the identification of relations between scientific terms in the natural language processing (NLP) community, within the context of human annotations on the abstracts of scholarly articles [10,20], seven relation types between scientific terms have been studied. They are HYPONYM-OF, PART-OF, USAGE, COMPARE, CONJUNCTION, FEATURE-OF, and RESULT. The annotations are in the form of generalized relation triples: ⟨experiment⟩ COMPARE ⟨another experiment⟩; ⟨method⟩ USAGE ⟨data⟩; ⟨method⟩ USAGE ⟨research task⟩. Since human language exhibits the paraphrasing phenomenon, identifying each specific relation between scientific concepts is impractical. In the framework of an automated pipeline for generating knowledge graphs over massive volumes of scholarly records, the task of classifying scientific relations (i.e., identify the appropriate relation type for each related concept pair from a set of predefined relations) is therefore indispensable.

In other text mining initiatives, comprehensive knowledge mining themes are being defined on scholarly investigations. The recent NLPContributionGraph Shared Task [19,17,16] released KG annotations of contributions including the facets of research problem, approach, experimental settings, and results, in an evaluation series that showed it a challenging task. Similarly, in the Life Sciences, comprehensive KGs from reports of biological assays, wet lab protocols and inorganic materials synthesis reactions and procedures [7,8,31,32,33,41] are released as ontologized machine-interpretable formats for training machine readers.

Digitalization based on Table Mining In the earlier subsection, we discussed information extraction models defined for retrieving the relevant structured information from the textual body of articles. Recent efforts are geared to mining

information from the semi-structured format of information in articles as tables. Unlike the high performances seen in information mining systems applied to textual data, text mining performances over tables are relatively much lower.

Milosevic et al. [39] tested methods for extracting numerical (number of patients, age, gender distribution) and textual (adverse reactions) information from tables identified by the `< table >` tag in the clinical literature as XML articles. Further, another line of work examined the classification of tables from HTML pages as entity, relational, matrix, list, and nondata leveraging specialized table embeddings called TabVec [21]. Wei et al.[46] defined a question answering task with data in Table cells as the answers over two different datasets, i.e. web data tables and news articles in text format tables. Another model called TAPAS [23] also addressed question answering over tabular data by extending BERT’s architecture to encode tables as input and training it end-to-end over tables crawled from Wikipedia. TableSeer [34] is a comprehensive tables mining search engine that crawls digital libraries, detects tables from documents, extracts their meta-data, and indexes and ranks tables in a user-friendly search interface.

Digitalization based on Textual and Tabulated Content Mining IBM’s science result extractor [24] first defined the *(Task, Dataset, Metric)* extraction task from articles. They trained a BERT classification model leveraging context data from the abstract, tables, and from table headers and captions. Their dataset comprised pdf-to-text converted articles from `PapersWithCode.com`. Following which, AxCell [30] presented an automated machine learning pipeline for extracting results from papers. It used several novel components, including table segmentation, to learn relevant structural knowledge to aid extraction. Unlike the first system, AxCell was trained and tested over \LaTeX source code of machine learning papers from `arXiv.org`. Furthermore, the SciRex corpus creation endeavor [26] defines mostly similar information targets as science result extractor. However, SciRex is evaluated on clean \LaTeX sources unlike the IBM extractor and our objective of trying to identify robust machine learning pipelines over articles in PDF format. Another recent system, SciNLP-KG [40], reformulated the LEADERBOARD extraction task as one with relation *evaluatedOn* between tasks and datasets, *evaluatedBy* between tasks and metrics, as well as coreferent and related relations between the same type of entities. Like us, they comprehensively investigated several transformer model variants. Nevertheless, owing to their task reformulation, relation-based evaluation, and different dataset our results cannot be compared. Finally, Hou et al., the developers of IBM’s science result extractor, recently released the TDMSci corpus [25] as a sequence labeling task for extracting *Task, Dataset, and Metric* at the sentence-level. We maintain the original document-level inference task definition.

In this paper, we investigate the science result extractor system [24], however, with a detailed empirical perspective. We comprehensively evaluate the potential of transformer models for the task by testing Bert, SciBert, and XLNet. We also test the models over a much larger empirical dataset emulating its application in practice in the framework of the scholarly digital libraries such as the ORKG.

3 Our Leaderboards Labeled Corpus

To facilitate supervised system development for the extraction of LEADERBOARDS from scholarly articles, we build an empirical corpus that encapsulates the task. LEADERBOARD extraction is essentially an inference task over the document. To alleviate the otherwise time-consuming and expensive corpus annotation task involving expert annotators, we leverage distant supervision from the available crowdsourced metadata in the PwC KB. In the remainder of this section, we explain our corpus creation and annotation process.

Scholarly Papers and Metadata from the PwC Knowledge Base. We created a new corpus as a collection of scholarly papers with their TDM triple annotations for evaluating the LEADERBOARDS extraction task inspired by the original IBM science result extractor [24] corpus. The collection of scholarly articles for defining our LEADERBOARD extraction objective is obtained from the publicly available crowdsourced leaderboards PapersWithCode (PwC)⁶. It predominantly represents articles in the Natural Language Processing and Computer Vision domains, among other AI domains such as Robotics, Graphs, Reasoning, etc. Thus, the corpus is representative for empirical AI research. The original downloaded collection (timestamp 2021-05-10 at 12:30:21)⁷ was pre-processed to be ready for analysis. While we use the same method here as the science result extractor, our corpus is different in terms of both labels and size, i.e. number of papers, as many more LEADERBOARDS have been crowdsourced and added to PwC since the original work.

PDF pre-processing While the respective articles’ metadata in machine-readable form was directly obtained from the PwC data release, the document itself being in PDF format needed to undergo pre-processing for pdf-to-text conversion so that its contents could be mined. For this, the GROBID parser [35] was applied to extract the title, abstract, and for each section, the section title and its corresponding content from the respective PDF article files. Each article’s parsed text was then annotated with TDM triples via distant labeling to create the final corpus.

Paper Annotation via Distant Labeling Each paper was associated with its LEADERBOARD TDM triple annotations. These were available as the crowdsourced metadata of each article in the PwC knowledge base (KB). The number of triples per article varied between 1 (minimum) and 54 (maximum) at an average of 4.1 labels per paper. The corpus was thus annotated as a distant labeling

⁶ <https://paperswithcode.com/>

⁷ Our corpus was downloaded from the PwC Github repository <https://github.com/paperswithcode/paperswithcode-data> and was constructed by combining the information in the files *All papers with abstracts* and *Evaluation tables* which included article urls and TDM crowdsourced annotation metadata.

	Ours		Original	
	Train	Test	Train	Test
Papers	3,753	1,608	170	167
“unknown” annotations	922	380	46	45
Total TDM-triples	11,724	5,060	327	294
Avg. number of TDM-triples per paper	4.1	4.1	2.64	2.41
Distinct TDM-triples	1,806	1,548	78	78
Distinct <i>Tasks</i>	288	252	18	18
Distinct <i>Datasets</i>	908	798	44	44
Distinct <i>Metrics</i>	550	469	31	31

Table 1. Ours vs. the original science result extractor [24] corpora statistics. The “unknown” labels were assigned to papers with no TDM-triples after the label filtering stage.

task since the labels for each paper were directly imported from the PwC KB without additional human curation of the varying forms of label names. Additionally, LEADERBOARDS that appeared in less than five papers were ignored. Consequently to the TDM labels filtering stage, some articles were without TDM triples and these articles were annotated with the label “unknown.”

Our overall corpus statistics are shown in Table 1. We adopted the 70/30 split for the Train/Test folds for the empirical system development (described in detail in section 6). In all, our corpus contained 5,361 articles split into 3,753 in the training data and 1,608 in the test data. There were unique TDM-triples overall. Note that since the test labels were a subset of the training labels, the unique labels overall can be considered as those in the training data. Table 1 also shows the distinct *Tasks*, *Datasets*, *Metrics* in the last three rows. Our corpus contains 288 *Tasks* defined on 908 *Datasets* and evaluated by 550 *Metrics*. This is significantly larger than the original corpus which had 18 *Tasks* defined on 44 *Datasets* and evaluated by 31 *Metrics*.

4 Leaderboard Extraction Task Definition

The task is defined on the dataset described in previous section. The dataset can be formalized as follows. Let p be a paper in the collection P . Each p is annotated with at least one triple (t_i, d_j, m_k) where t_i is the i^{th} task defined, d_j the j^{th} dataset and m_k the k^{th} system evaluation metric. The number of triples per paper vary.

In the supervised inference task, the input data instance corresponds to the pair: a paper p represented as the DocTAET context feature $p_{DocTAET}$ and its TDM-triple (t, d, m) . The inference data instance, then is $(c; [(t, d, m), p_{DocTAET}])$ where $c \in \{true, false\}$ is the inference label. Thus, specifically, our LEADERBOARD extraction problem is formulated as a *natural language inference* task between the DocTAET context feature $p_{DocTAET}$ and the (t, d, m) triple annotation. (t, d, m) is *true* if it is among the paper’s TDM-triples, otherwise *false*. The *false* instances are artificially created by random selection of (t, d, m) annotations from another paper. Cumulatively, LEADERBOARD construction is a multi-label, multi-class inference problem.

DocTAET Context Representation Feature

Title:

Deep Recurrent Generative Decoder for Abstractive Text **Summarization** *

Abstract:

We propose a new framework for abstractive text summarization based on a sequence-to-sequence oriented encoder-decoder model equipped with a deep recurrent generative decoder (DRGN). Latent structure information implied in the target summaries is learned based on a recurrent latent random model for improving the **summarization** quality. Neural variational inference is employed to address the intractable posterior inference for the recurrent latent variables. Abstractive summaries are generated based on both the generative latent variables and the discriminative deterministic states. Extensive experiments on some benchmark datasets in different languages show that DRGN achieves improvements over the state-of-the-art methods.

Experimental Setup:

We use **ROUGE** (score) as our evaluation metric with standard options For the experiments on the English dataset **Gigawords**, we set the dimension of word embeddings to 300, and the dimension of hidden states and latent variables to 500 For the dataset of LCSTS, the dimension of word embeddings is 350 The comparison results on the validation datasets of **Gigawords** and LCSTS are shown in Actually, the performance of the standard The results on the English datasets of **Gigawords** and **DUC-2004** are shown in and respectively In fact, extracting all such features is a time consuming work, especially on large-scale datasets such as **Gigawords** The results on the Chinese dataset LCSTS are shown in

Table Info:

Table 1: **ROUGE** - F1 on validation sets R - 1 R - 2 R - L Table 1 : **ROUGE** - F1 on validation sets Table 2: **ROUGE**-F1 on **Gigawords** Table 3 : **ROUGE** - Recall on **DUC2004** R - 1 R - 2 R - L Table 2 : **ROUGE** - F1 on **Gigawords** Table 3: **ROUGE**-Recall on **DUC2004** Table 4 : **ROUGE** - F1 on LCSTS Table 3 : **ROUGE** - Recall on **DUC2004** R - 1 R - 2 R - L Table 2 : **ROUGE** - F1 on **Gigawords** Table 4: **ROUGE**-F1 on LCSTS Table 4 : **ROUGE** - F1 on LCSTS R - 1 R - 2 Table 3 : **ROUGE** - Recall on **DUC2004** R - L

Leaderboard triples

(**Summarization**; **DUC 2004 Task 1**; **ROUGE-L**)
(**Summarization**; **Gigaword**; **ROUGE-1**)
(**Summarization**; **DUC 2004 Task 1**; **ROUGE-1**)
(**Summarization**; **DUC 2004 Task 1**; **ROUGE-2**)
(**Summarization**; **Gigaword**; **ROUGE-L**)
(**Summarization**; **Gigaword**; **ROUGE-2**)

Fig. 1. The DocTAET model with context features as a concatenation of the Scholarly **Document's** **T**itle, **A**bstract, **E**xperimental-Setup, and **T**able content/captions for training NLI transformer models on a set of LEADERBOARD triples. The figure illustrates specifically six (*task*, *dataset*, *metric*) triples and their context in the original article text extracted as the feature for transformer models.

4.1 DocTAET Context Feature

In Figure 1, we depict the DocTAET context feature [24]. Essentially, the LEADERBOARD extraction task is defined on the full document content. However, the respective (*task*, *dataset*, *metric*) label annotations are mentioned only in specific places in the full paper such as in the Title, Abstract, Introduction, Tables. The DocTAET feature was thus defined to capture the targeted context information to facilitate the (*task*, *dataset*, *metric*) triple inference. It focused on capturing the context from four specific places in the text-parsed article, i.e. from the title, abstract, first few lines of the experimental setup section as well as table content and captions.

5 Transformer-based Leaderboard Extraction Models

For LEADERBOARD extraction [24], we employ deep transfer learning modeling architectures that rely on a recently popularized neural architecture – the transformer [44]. Transformers are arguably the most important architecture for natural language processing (NLP) today since they have shown and continue to show impressive results in several NLP tasks [15]. Owing to the self-attention mechanism in the transformer models, they can be fine-tuned on many downstream tasks. These models have thus crucially popularized the transfer learning paradigm in NLP. We investigate three transformer-based model variants for LEADERBOARD extraction in a Natural Language Inference configuration.

Natural language inference (NLI), generally, is the task of determining whether a “hypothesis” is true (entailment), false (contradiction), or undetermined (neutral) given a “premise” [2]. For LEADERBOARD extraction, the slightly adapted NLI task is to determine that the *(task, dataset, metric)* “hypothesis” is true (entailed) or false (not entailed) for a paper given the “premise” as the DocTAET context feature representation of the paper.

Currently, there exist several transformer-based models. In our experiments, we investigated three core models: two variants of Bert, i.e. the vanilla Bert [15] and the scientific Bert (SciBert) [11]. We also tried a different type of transformer model than Bert called XLNet [47] which employs Transformer-XL as the backbone model. Next, we briefly describe the three variants we use.

5.1 Bert Models

Bert (i.e., Bidirectional Encoder Representations from Transformers), is a bidirectional autoencoder (AE) language model. As a pre-trained language representation built on the deep neural technology of transformers, it provides NLP practitioners with high-quality language features from text data simply out-of-the-box and thus improves performance on many NLP tasks. These models return contextualized word embeddings that can be directly employed as features for downstream tasks [28].

The first Bert model we employ is Bert_{base} (12 layers, 12 attention heads, and 110 million parameters) which was pre-trained on billions of words from the BooksCorpus (800M words) and the English Wikipedia (2,500M words).

The second Bert model we employ is the pre-trained scientific Bert called SciBert [11]. SciBert was pretrained on a large corpus of scientific text. In particular, the pre-training corpus is a random sample of 1.14M papers from Semantic Scholar⁸ consisting of full texts of 18% of the papers from the computer science domain and 82% from the broad biomedical domain. For both Bert_{base} and SciBert, we used their uncased variants.

⁸ <https://semanticscholar.org>

5.2 XLNet

XLNet is an autoregressive (AR) language model [47] that enables learning bidirectional contexts using Permutation Language Modeling, unlike Bert which uses Masked Language Modeling. Thus in PLM all tokens are predicted but in random order, whereas in MLM only the masked (15%) tokens are predicted. This is also in contrast to the traditional language models, where all tokens were predicted in sequential order instead of random order. Random order prediction helps the model to learn bidirectional relationships and therefore better handle dependencies and relations between words. In addition, it uses Transformer XL [14] as the base architecture, which models long contexts unlike the Bert models with contexts limited to 512 tokens. Since only cased models are available for XLNet, we used the cased XLNet_{base} (12 layers, 12 attention heads, and 110 million parameters).

6 Automated Leaderboard Mining

6.1 Experimental Setup

Parameter Tuning. We used the Hugging Transformer libraries⁹ with their Bert variants and XLNet implementations. In addition to the standard fine-tuned setup for NLI, the transformer models were trained with a learning rate of $1e - 5$ for 14 epochs; and used the *AdamW* optimizer with a weight decay of 0 for *bias*, *gamma*, *beta* and 0.01 for the others. Our models’ hyperparameters details are available online.¹⁰

In addition, we introduce a task-specific parameter that was crucial in obtaining optimal task performance from the models. It was the number of *false* triples per paper. This parameter controls the discriminatory ability of the model. The original science result extractor system [24] considered $|n| - |t|$ *false* instances for each paper, where $|n|$ was the distinct set of triples overall and $|t|$ was the number of *true* LEADERBOARD triples per paper. This approach would not generalize to our larger corpus with over 2,500 distinct triples. In other words, considering that each paper had on average 4 *true* triples, it would have 2,495 *false* triples which would strongly bias the classifier learning toward only *false* inferences. Thus, we tuned this parameter in a range of values in the set {10, 50, 100} which at each experiment run was fixed for all papers.

Finally, we imposed an artificial trimming of the DocTAET feature to accommodate the Bert models maximum token length of 512. For this, the token lengths of the experimental setup and table info were initially truncated to approximately 150, after which the complete DocTAET feature is trimmed to 512 tokens.

⁹ <https://github.com/huggingface/transformers>

¹⁰ https://github.com/Kabongosalomon/task-dataset-metric-nli-extraction/blob/main/train_tdm.py

	Macro P	Macro R	Macro F1	Micro P	Micro R	Micro F1
(a) Task + Dataset + Metric Extraction						
TDM-IE _{Bert}	62.5	75.2	65.3	60.8	76.8	67.8
ORKG-TDM _{Bert}	68.1	67.5	65.5	79.6	63.3	70.5
ORKG-TDM _{SciBert}	65.7	77.2	68.3	65.7	76.8	70.8
ORKG-TDM _{XLNet}	71.7	73.9	70.6	77.1	70.9	73.9
(b) Task + Dataset + Metric Extraction (without "Unknown" annotation)						
TDM-IE _{Bert}	54.1	65.9	56.6	60.2	73.1	66.0
ORKG-TDM _{Bert}	59.0	55.4	54.7	79.5	57.6	66.8
ORKG-TDM _{SciBert}	57.6	68.7	60.1	65.3	73.1	69.0
ORKG-TDM _{XLNet}	63.5	64.1	61.4	76.4	66.4	71.1

Table 2. LEADERBOARD triple extraction task, comparison of our ORKG-TDM models versus the original science result extractor model (first row in parts (a) and (b)) on the original corpus (see last two columns in Table 1 for the original corpus statistics).

Two-Fold Cross Validation. To evaluate robust models, we performed two-fold cross validation experiments. In each fold experiment, we train a model on 70% of the overall dataset, and test on the remaining 30% ensuring that the test data splits are not identical between the folds. Thus, all cumulative results reported are averaged results over the two folds. Also, Table 1 corpus statistics are averaged estimates over the two experimental folds.

Evaluation Metrics. Within the two-fold experimental settings, we report macro- and micro-averaged precision, recall, and F1 scores for our LEADERBOARD extraction task on the test dataset. The macro scores capture the averaged class-level task evaluations, whereas the micro scores represent fine-grained instance-level task evaluations.

Further, the macro and micro evaluation metrics for the overall task have two evaluation settings: 1) considers papers with LEADERBOARDs and papers with “unknown” in the metric computations; 2) only papers with LEADERBOARDs are considered while the papers with “unknown” are excluded.

6.2 Experimental Results

The results from our comprehensive evaluations with respect to four main research questions (noted as **RQ1**, **RQ2**, **RQ3**, and **RQ4**) are shown in Tables 2, 3, and 4, respectively.

RQ1: How well do our transformer models perform for Leaderboard extraction compared to the original science result extractor when trained in the identical original experimental setting? In the last two columns in Table 1, we showed statistics of the comparatively smaller original corpus that defined and evaluated the LEADERBOARD extraction task [24]. As a quick recap, the original corpus had 78 distinct TDM-triples including “unknown,” and a distribution of 170 papers in the train dataset and 167 papers in the test dataset as fixed partitions; there were 46 and 45 “unknown” papers in the train and test sets, respectively. We evaluate all the three transformer models

	Macro P	Macro R	Macro F1	Micro P	Micro R	Micro F1
Average Evaluation Accross 2-fold						
ORKG-TDM _{Bert}	92.8	93.9	92.4	95.5	89.1	92.1
ORKG-TDM _{SciBert}	90.9	93.4	91.1	94.1	88.5	91.2
ORKG-TDM _{XLNet}	92.8	94.8	92.8	94.9	91.2	93.0
Average Evaluation Accross 2-fold (without "Unknown" annotation)						
ORKG-TDM _{Bert}	91.7	92.1	90.8	95.7	88.3	91.8
ORKG-TDM _{SciBert}	89.7	91.4	89.4	94.4	87.6	90.9
ORKG-TDM _{XLNet}	91.6	93.1	91.2	95.0	90.5	92.7

Table 3. Top results for Bert (10neg; 1,302unk), SciBert (10neg; 1,302unk), and XLNet (10neg; 1,302unk)

Entity	Macro			Micro		
	P	R	F ₁	P	R	F ₁
TDM	91.6	93.1	91.2	95.0	90.5	92.7
Task	93.7	94.8	93.6	97.4	93.6	95.5
Dataset	92.9	93.6	92.4	96.6	91.5	94.0
Metric	92.5	94.2	92.5	96.0	92.5	94.2

Table 4. Performance of our best model, i.e. ORKG-TDM_{XLNet}, for *Task*, *Dataset*, and *Metric* concept extraction of the LEADERBOARD

Document Representation	Macro P	Macro R	Macro F1	Micro P	Micro R	Micro F1
Title + Abstract	88.6	92.9	89.4	92.6	90	91.3
Title + Abstract + ExpSetup	89.2	91.5	89.2	94.2	89	91.5
Title + Abstract + TableInfo	90.5	94.4	91.2	93.5	93.2	93.3
Title + Abstract + ExpSetup + TableInfo	92.3	93.5	91.7	95.1	92	93.5

Table 5. Ablation results of our best model, i.e. ORKG-TDM_{XLNet}, for LEADERBOARD extraction as (*task*, *dataset*, *metric*) triples

on this original corpus to compare our model performances. These results are shown in Table 2. As we can see in the table, all three of our models outperform the original with XLNet reporting the best score. We obtain a 70.6 macro F1 versus 65.3 in the baseline. The Bert model is only a few fractional points better at 65.5. Further, we obtain a 73.9 micro F1 versus 67.8 in the baseline, with the Bert model at 70.5. With these results our models outperform the original system in the original settings reported for this task.

Next, we examine the results of our models on our larger corpus for LEADERBOARD extraction.

RQ2: How do the transformer models perform on a large corpus for Leaderboard construction? We examine this question in light of the results reported in Table 3. We find that again, consistent with the observations made in Table 2, XLNet outperforms Bert and SciBert. Note that we have leveraged XLNet with the limited context length of 512 tokens (by truncating parts of the context described in Section 6.1) and thus the potential of leveraging the full context in XLNet models remains untapped. Nevertheless, XLNet still distinguishes itself from the Bert models with its PLM language modeling objective versus Bert’s MLM. This in practice has also shown to perform better [47]. Thus XLNet with limited context encoding can still outperform the Bert models.

Among all the models, we find SciBert shows a slightly lower performance compared to Bert. This is a slight variation on performance observations obtained in the original smaller dataset (results in Table 2) where Bert was slightly lower than SciBert. These results indicate that in smaller datasets, the SciBert model should be preferred since its underlying science-specific pretrained corpus would compensate for signal absences in the task training corpus. However, with larger training datasets for finetuning the Bert models, the underlying pretraining corpus domains show not to be critical to the overall model performance.

Further, we observe the macro and micro evaluations for all three systems (ORKG-TDM_{Bert}, ORKG-TDM_{SciBert}, ORKG-TDM_{XLNet}) show evenly balanced, similar scores. This tells us that our models handles the majority and minority TDM classes evenly. Thus given our large training corpus, the transformers remain unaffected by the underlying dataset class distributions.

RQ3: Which of the three Leaderboard concepts are easy or challenging to extract? As a fine-grained examination of our best model, i.e. ORKG-TDM_{XLNet}, we examined its performance for extracting each of three concepts (*Task, Dataset, Metric*) separately. These results are shown in Table 4. From the results, we observe that TASK is the easiest concept to extract, followed by METRIC, and then DATASET. We ascribe the low performance for extracting the DATASET concept due to the variability in its naming seen across papers even when referring to the same real-world entity. For example, the real-world dataset entity *CIFAR-10* is labeled as *CIFAR-10, 4000 Labels* in some papers and *CIFAR-10, 250 Labels* in others. This phenomenon is less prevalent for TASK and the METRIC concepts. For example, the TASK ‘Question Answering’ is rarely referenced differently across papers addressing the task. Similarly, for METRIC, ACCURACY, as an example, has very few variables namings.

RQ4: Which aspect of the DocTAET context feature representation had the highest impact for Leaderboard extraction? Further, in Table 5, we breakdown performance of our best model, i.e. ORKG-TDM_{XLNet}, examining the impact of the features as the shortened context from the articles for TDM inference. We observe, that adding additional contextual information in addition to title and abstract increases the performance significantly, while the actual type of additional information (i.e. experimental setup, table information or both) impacts the performance to a lower extend.

6.3 Integrating ORKG-TDM in Scholarly Digital Libraries

Ultimately, with LEADERBOARD construction we aim to give researchers oversight over the state-of-the-art wrt. certain research questions (i.e. tasks). Figure 2 shows how our ORKG-TDM LEADERBOARD mining is integrated into the Open Research Knowledge Graph scholarly knowledge platform. The fact, that the automated LEADERBOARD mining results in incomplete and partially imprecise results can be alleviated by the crowdsourcing and curation features implemented

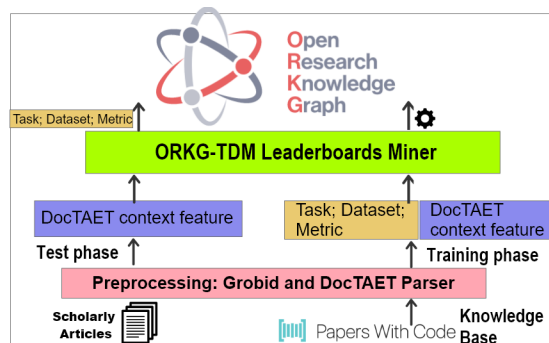


Fig. 2. Overview of the LEADERBOARD Extraction Machine Learning Process Flow for ORKG-TDM in the Open Research Knowledge Graph Digital Library

in the ORKG. Also, ORKG provides features for dynamic LEADERBOARD visualization, publication, versioning etc.

Although, the experiments of our study targeted empirical AI research, we are confident, that the approach is transferable to similar scholarly knowledge extraction tasks in other domains. For example in chemistry or material sciences, experimentally observed properties of substances or materials under certain conditions could be obtained from various papers.

7 Conclusion and Future Work

In this paper, we investigated the LEADERBOARD extraction task w.r.t. three different transformer-based models. Our overarching aim with this work is to build a system for comparable scientific concept extractors from scholarly articles. Therefore as a next step, we will extend the current triples (task, dataset, metric) model with additional concepts that are suitable candidates for a LEADERBOARD such as *score* or *code urls*, etc. In this respect, we will adopt a hybrid system wherein some elements will be extracted by the machine learning system as discussed in this work while other elements will be extracted by a system of rules and regular expressions. Also, we plan to combine the automated techniques presented herein with a crowdsourcing approach for further validating the extracted results and providing additional training data. Our work in this regard is embedded in a larger research and service development agenda, where we build a comprehensive knowledge graph for representing and tracking scholarly advancements [27]. We also envision the task-dataset-metric extraction approach to be transferable to other domains (such as materials science, engineering simulations etc.). Our ultimate target is to create a comprehensive structured knowledge graph tracking scientific progress in various scientific domains, which can be leveraged for novel machine-assistance measures in scholarly communication, such as question answering, faceted exploration and contribution correlation tracing.

Acknowledgements This work was co-funded by the Federal Ministry of Education and Research (BMBF) of Germany for the project LeibnizKILabor (grant no. 01DD20003) and by the European Research Council for the project ScienceGRAPH (Grant agreement ID: 819536).

References

1. Ai metrics. <https://www.eff.org/ai/metrics>, accessed: 2021-04-26
2. Natural Language Inference. <https://paperswithcode.com/task/natural-language-inference>, accessed: 22 April 2021
3. Nlp-progress. <http://nlpprogress.com/>, accessed: 2021-04-26
4. paperswithcode.com. <https://paperswithcode.com/>, accessed: 2021-04-26
5. Reddit sota. <https://github.com/RedditSota/state-of-the-art-result-for-machine-learning-problems>, accessed: 2021-04-26
6. Squad explorer. <https://rajpurkar.github.io/SQuAD-explorer/>, accessed: 2021-04-26
7. Anteghini, M., D’Souza, J., Dos Santos, V.A.M., Auer, S.: Scibert-based semantification of bioassays in the open research knowledge graph. In: EKAW-PD 2020. pp. 22–30 (2020)
8. Anteghini, M., D’Souza, J., dos Santos, V.A.M., Auer, S.: Representing semantified biological assays in the open research knowledge graph. In: International Conference on Asian Digital Libraries. pp. 89–98. Springer (2020)
9. Auer, S.: Towards an open research knowledge graph (Jan 2018). <https://doi.org/10.5281/zenodo.1157185>, <https://doi.org/10.5281/zenodo.1157185>
10. Augenstein, I., Das, M., Riedel, S., Vikraman, L., McCallum, A.: Semeval 2017 task 10: Scienceie - extracting keyphrases and relations from scientific publications. In: SemEval@ACL (2017)
11. Beltagy, I., Lo, K., Cohan, A.: Scibert: A pretrained language model for scientific text. arXiv preprint arXiv:1903.10676 (2019)
12. Brack, A., D’Souza, J., Hoppe, A., Auer, S., Ewerth, R.: Domain-independent extraction of scientific concepts from research articles. In: European Conference on Information Retrieval. pp. 251–266. Springer (2020)
13. Chiarelli, A., Johnson, R., Richens, E., Pinfield, S.: Accelerating scholarly communication: The transformative role of preprints (2019)
14. Dai, Z., Yang, Z., Yang, Y., Carbonell, J.G., Le, Q., Salakhutdinov, R.: Transformer-xl: Attentive language models beyond a fixed-length context. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. pp. 2978–2988 (2019)
15. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
16. D’Souza, J., Auer, S., Pedersen, T.: SemEval-2021 task 11: Nlpcontributiongraph - structuring scholarly nlp contributions for a research knowledge graph. In: Proceedings of the Fifteenth Workshop on Semantic Evaluation. Association for Computational Linguistics, Bangkok (online) (Aug 2021)
17. D’Souza, J., Auer, S., Pederson, T.: SemEval-2021 Task 11: NLPContributionGraph - Structuring Scholarly NLP Contributions for a Research Knowledge Graph (May 2021), <https://zenodo.org/record/4737071>

18. D'Souza, J., Hoppe, A., Brack, A., Jaradeh, M.Y., Auer, S., Ewerth, R.: The stem-ecr dataset: Grounding scientific entity references in stem scholarly content to authoritative encyclopedic and lexicographic sources. In: LREC. pp. 2192–2203. Marseille, France (May 2020)
19. D'Souza, J., Auer, S.: Sentence, phrase, and triple annotations to build a knowledge graph of natural language processing contributions—a trial dataset. *Journal of Data and Information Science* p. 20210429 (2021)
20. Gábor, K., Buscaldi, D., Schumann, A.K., QasemiZadeh, B., Zargayouna, H., Charnois, T.: Semeval-2018 task 7: Semantic relation extraction and classification in scientific papers. In: Proceedings of The 12th International Workshop on Semantic Evaluation. pp. 679–688 (2018)
21. Ghasemi-Gol, M., Szekely, P.: Tabvec: Table vectors for classification of web tables. arXiv preprint arXiv:1802.06290 (2018)
22. Handschuh, S., QasemiZadeh, B.: The acl rd-tec: a dataset for benchmarking terminology extraction and classification in computational linguistics. In: COLING 2014: 4th international workshop on computational terminology (2014)
23. Herzig, J., Nowak, P.K., Mueller, T., Piccinno, F., Eisenschlos, J.: Tapas: Weakly supervised table parsing via pre-training. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 4320–4333 (2020)
24. Hou, Y., Jochim, C., Gleize, M., Bonin, F., Ganguly, D.: Identification of tasks, datasets, evaluation metrics, and numeric scores for scientific leaderboards construction. arXiv preprint arXiv:1906.09317 (2019)
25. Hou, Y., Jochim, C., Gleize, M., Bonin, F., Ganguly, D.: Tdmsci: A specialized corpus for scientific literature entity tagging of tasks datasets and metrics. In: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume. pp. 707–714 (2021)
26. Jain, S., van Zuylen, M., Hajishirzi, H., Beltagy, I.: Scirex: A challenge dataset for document-level information extraction. In: Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. pp. 7506–7516 (2020)
27. Jaradeh, M.Y., Oelen, A., Farfar, K.E., Prinz, M., D'Souza, J., Kismihók, G., Stocker, M., Auer, S.: Open research knowledge graph: next generation infrastructure for semantic scholarly knowledge. In: Proceedings of the 10th International Conference on Knowledge Capture. pp. 243–246 (2019)
28. Jiang, M., D'Souza, J., Auer, S., Downie, J.S.: Improving scholarly knowledge representation: Evaluating bert-based models for scientific relation classification. In: International Conference on Asian Digital Libraries. pp. 3–19. Springer (2020)
29. Jinha, A.E.: Article 50 million: an estimate of the number of scholarly articles in existence. *Learned Publishing* **23**(3), 258–263 (2010)
30. Kardas, M., Czapla, P., Stenetorp, P., Ruder, S., Riedel, S., Taylor, R., Stojnic, R.: Axccl: Automatic extraction of results from machine learning papers. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP). pp. 8580–8594 (2020)
31. Kononova, O., Huo, H., He, T., Rong, Z., Botari, T., Sun, W., Tshitoyan, V., Ceder, G.: Text-mined dataset of inorganic materials synthesis recipes. *Scientific data* **6**(1), 1–11 (2019)
32. Kulkarni, C., Xu, W., Ritter, A., Machiraju, R.: An annotated corpus for machine reading of instructions in wet lab protocols. In: NAACL: HLT, Volume 2 (Short Papers). pp. 97–106. New Orleans, Louisiana (Jun 2018). <https://doi.org/10.18653/v1/N18-2016>

33. Kuniyoshi, F., Makino, K., Ozawa, J., Miwa, M.: Annotating and extracting synthesis process of all-solid-state batteries from scientific literature. In: LREC. pp. 1941–1950 (2020)
34. Liu, Y., Bai, K., Mitra, P., Giles, C.L.: Tableseer: automatic table metadata extraction and searching in digital libraries. In: Proceedings of the 7th ACM/IEEE-CS joint conference on Digital libraries. pp. 91–100 (2007)
35. Lopez, P.: Grobid: Combining automatic bibliographic data recognition and term extraction for scholarship publications. In: International conference on theory and practice of digital libraries. pp. 473–474. Springer (2009)
36. Luan, Y., He, L., Ostendorf, M., Hajishirzi, H.: Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. arXiv preprint arXiv:1808.09602 (2018)
37. Luan, Y., He, L., Ostendorf, M., Hajishirzi, H.: Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction. In: EMNLP (2018)
38. Manning, C.D.: Computational linguistics and deep learning. *Computational Linguistics* **41**(4), 701–707 (2015)
39. Milosevic, N., Gregson, C., Hernandez, R., Nenadic, G.: A framework for information extraction from tables in biomedical literature. *Int. J. Doc. Anal. Recognit.* **22**(1), 55–78 (Mar 2019). <https://doi.org/10.1007/s10032-019-00317-0>, <https://doi.org/10.1007/s10032-019-00317-0>
40. Mondal, I., Hou, Y., Jochim, C.: End-to-end nlp knowledge graph construction. arXiv preprint arXiv:2106.01167 (2021)
41. Mysore, S., Jensen, Z., Kim, E., Huang, K., Chang, H.S., Strubell, E., Flanigan, J., McCallum, A., Olivetti, E.: The materials science procedural text corpus: Annotating materials synthesis procedures with shallow semantic structures. In: Proceedings of the 13th Linguistic Annotation Workshop. pp. 56–64 (2019)
42. Oelen, A., Stocker, M., Auer, S.: Crowdsourcing scholarly discourse annotations. In: 26th International Conference on Intelligent User Interfaces. pp. 464–474 (2021)
43. Renear, A.H., Palmer, C.L.: Strategic reading, ontologies, and the future of scientific publishing. *Science* **325**(5942), 828–832 (2009)
44. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in neural information processing systems. pp. 5998–6008 (2017)
45. Ware, M., Mabe, M.: The stm report: An overview of scientific and scholarly journal publishing (03 2015)
46. Wei, X., Croft, B., Mccallum, A.: Table extraction for answer retrieval. *Inf. Retr.* **9**(5), 589–611 (Nov 2006). <https://doi.org/10.1007/s10791-006-9005-5>, <https://doi.org/10.1007/s10791-006-9005-5>
47. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R., Le, Q.V.: Xlnet: Generalized autoregressive pretraining for language understanding. arXiv preprint arXiv:1906.08237 (2019)

A Examples of Leaderboards in Our Corpus

The top-3 most common LEADERBOARDS in our training set included: 1) *Image Classification, ImageNet, Top 1 Accuracy*; 2) *Object Detection, COCO test-dev, box AP*; 3) *Image Classification, CIFAR-10, Percentage correct* occurring 93, 57, and 51 times, respectively.

The top-3 least common randomly selected LEADERBOARDS in our training set included: 1) *Word Sense Disambiguation, WiC-TSV, Task 1 Accuracy: all*; 2) *Entity Linking, WiC-TSV, Task 1 Accuracy: all*; 3) *Causal Inference, IDHP, Average Treatment Effect Error* occurring once.