

# Joint Representations of Text and Knowledge Graphs for Retrieval and Evaluation

Teven Le Scao\*

Université de Lorraine, Hugging Face  
teven.le-scao@loria.fr

Claire Gardent

CNRS/LORIA, Université de Lorraine  
claire.gardent@loria.fr

## Abstract

A key feature of neural models is that they can produce semantic vector representations of objects (texts, images, speech, etc.) ensuring that similar objects are close to each other in the vector space. While much work has focused on learning representations for other modalities, there are no aligned cross-modal representations for text and knowledge base (KB) elements. One challenge for learning such representations is the lack of parallel data, which we use contrastive training on heuristics-based datasets and data augmentation to overcome, training embedding models on (KB graph, text) pairs. On WEBNLG, a cleaner manually crafted dataset, we show that they learn aligned representations suitable for retrieval. We then fine-tune on annotated data to create EREDAT (Ensembled Representations for Evaluation of DATA-to-Text), a similarity metric between English text and KB graphs. EREDAT outperforms or matches state-of-the-art metrics in terms of correlation with human judgments on WEBNLG even though, unlike them, it does not require a reference text to compare against.

## 1 Introduction

Neural approaches have progressed in capturing semantic relatedness between larger and larger text units, from Word2Vec (Mikolov et al., 2013) to SBERT (Reimers and Gurevych, 2019). Such models have shown to perform well on a wide array of semantic similarity tasks, helped in part by retrieval systems like DPR (Karpukhin et al., 2020a).

Other work has shown that deep representations of knowledge bases (KBs) help improve such tasks as few shot link prediction, analogical reasoning (Pezeshkpour et al., 2018; Pahuja et al., 2021), entity linking (Yu et al., 2020) or cross-lingual entity alignment (Chen et al., 2018; Xu et al., 2019).

In this work, we focus on learning cross-modal representations for English text and KB graphs. Our input graphs are in RDF (Resource Description Framework, (Miller, 1998)) format, a standard where graphs are sets of (*subject, predicate, object*) triples. We linearize those graphs and consider them as text data so that the same model can take text and graphs as input. Given some aligned RDF-text data, our model learns fixed-length latent representations for texts and RDF graphs such that texts and RDF graphs that are semantically similar are close in vector space. This enables retrieval across modalities and allows us to create a cross-modality similarity score which can be used to evaluate the output of RDF-to-text generation models.

One challenge for learning cross-modal RDF-text representations is the lack of parallel data. We train on various RDF-text datasets created using distant supervision techniques, either combining these datasets or using them in isolation. We then compare the performance of the resulting retrieval models (i) on the WEBNLG dataset, a parallel RDF-text dataset where texts are crowdsourced to match the graph (texts and graphs are semantically equivalent), and (ii) on WIKICHUNKS, a more challenging, less well aligned dataset which imitates the conditions in which retrieval on Wikipedia is usually executed. We use the difference in performance between models to analyze the alignment quality of training datasets.

Distance within embedding space can be used to evaluate the output of RDF-to-text generation models (Is the generated text similar to the input graph?). In order to evaluate this metric, we compute correlations between our model’s similarity score for graph-text pairs and human judgments of semantic adequacy (input/output semantic similarity) using ratings from the 2020 WEBNLG Challenge. After fine-tuning on data from the 2017 WEBNLG challenge, as well as introducing new classes of data augmentation at pre-training time,

\*Corresponding author. teven.lescao@gmail.com

our best system, EREDAT, is better or on par than existing metrics at correlating with human evaluation, even though it does not require a reference for comparison as do most NLG evaluation metrics such as BLEU (Papineni et al., 2002), TER (Snover et al., 2006), BLEURT (Sellam et al., 2020b), METEOR (Banerjee and Lavie, 2005) or BERT-Score (Zhang\* et al., 2020).

Our contributions can be summarised as follows.

- We train a cross-modal RDF-text model to learn aligned (RDF graph, text) representations, making it suitable for cross-modal retrieval. We show that this retrieval model outperforms a state-of-the-art text-only retrieval model by a large margin, demonstrating the effectiveness of our adaptation procedure. We train on several datasets of RDF-text pairs, using the quality of the ensuing retrieval models to analyze the quality of training datasets.
- We provide a novel evaluation metric for RDF-to-text generation models by combining bi- and cross-encoder training procedures and adding adversarial data to address the models’ weaknesses. We show that this new metric outperforms other existing RDF-to-text evaluation metrics in terms of correlation with human judgments of semantic adequacy, even though it does not require a costly human reference to compare against. We release our models on [huggingface.co](https://huggingface.co) under the Apache 2.0 license.

## 2 Related Work

We briefly review recent approaches to uni- and cross-modal retrieval, representation learning models, and evaluation metrics for Natural Language Generation (NLG) models.

**Natural Language Retrieval Models.** For natural language, a first class of retrieval models focuses on retrieving sentences that are similar to some input sentence. BERT (Devlin et al., 2019) has been used as a cross-encoder. Two sentences are given with a separator token, cross-attention applies to all input tokens and the resulting representation is fed into a linear layer to score the match. However, this is computationally inefficient as it is not possible to pre-compute and index such representations. A pre-computable model was proposed by (Reimers and Gurevych, 2019) who

used twin encoders pre-trained on Natural Language Inference data (Bowman et al., 2015) to set new state-of-the-art performance on a large set of sentence scoring tasks. Further work (Chen et al., 2020; Humeau et al., 2019) combined cross- and bi-encoders to reach a tradeoff between accuracy and efficiency. We differ from those works in that we focus on cross-modal representation learning.

### **Representation Learning for Knowledge-Bases.**

Various KB embedding models have been proposed to support downstream applications such as KB completion or alignment of different bases. Compositional approaches (Nickel et al., 2011, 2016) use tensor products to model relations as functions of their argument entities. Translational approaches model relations as translation operations from the subject (head) to object (tail) entity (Bordes et al., 2013; Yang et al., 2014; Trouillon et al., 2016). Neural models have also leveraged 2-D convolutions over entity embeddings to predict relations (Dettmers et al., 2018) as well as graph convolutional networks (Schlichtkrull et al., 2018). All these approaches focus on representation learning for Knowledge-Bases entities and relations. In contrast, we focus on cross-modal similarity between a text and a KB graph.

### **Cross-Modal Representation Learning and Retrieval.**

Some work has focused on incorporating natural language information to improve KB representations. (Han et al., 2016; Toutanova et al., 2015; Wu et al., 2016) encode words and KB entities into a single vector space, and (Wang and Li, 2016; Yamada et al., 2016) learn word and entity embeddings separately then map them into a shared space. Both approaches use text as additional training signal to improve KB representations, and limit themselves to word-level information. Instead, we focus on scoring the similarity between arbitrary-length natural language text and a KB graph. We are not aware of any extant such text-KB models. The best-known cross-modal contrastive model is Radford et al. (2021), which pre-trained an image-text match scoring model.

### **Evaluation metrics for Natural Language Generation Models.**

Surface-based metrics such as BLEU (Papineni et al., 2002), which measure token overlap between generated and reference text, are commonly used. Methods such as BERT-Score (Zhang\* et al., 2020) or BLEURT (Sellam et al., 2020a) which leverage neural representations are

currently state-of-the-art. All these methods compute a score by comparing the generated text with human-produced references, rarely available and costly to produce. Some metrics evaluate the generated output with respect to the input rather than to a reference. Wiseman et al. (2017) use the precision of input relations found in the output texts. (Dušek and Kasner, 2020) use a natural language inference pre-trained model to score input-output two-way entailment. For data-to-text generation specifically, (Rebuffel et al., 2021) introduce Data-QuestEval, which uses question answering to compare input graph and output text.

### 3 Learning Cross-Modal RDF-text Representations

#### 3.1 Model

Similar to (Schroff et al., 2015; Reimers and Gurevych, 2019), we use twin Transformer encoders to create RDF and text representations such that the embeddings of an RDF graph and of a piece of text with similar content are close in the vector space. A mean-pooling operation creates fixed-sized embeddings  $embed(x)$  for  $x$  either an RDF graph or a text. RDF graphs are linearized as:

```
[S] <subject1> [P] <property1> [O]
<object1> ... [S] <subjectn> [P]
<propertyn> [O] <objectn>
```

where "[S]", "[P]", "[O]" serve as special tokens and are added to the tokenizer vocabulary. This allows us to treat any knowledge base format.

We train this system using a contrastive loss with *in-batch negatives* (Henderson et al., 2017). This variant of contrastive loss computes the pairwise similarities between every text and every RDF in the batch. A softmax is then applied on the RDF axis, which creates a multi-class classification problem: every text data point must be matched to the parallel RDF. The loss can be written as :

$$l = - \sum_{i \in I} \log \left( \frac{\exp(\text{sim}(\text{text}_i, \text{rdf}_i))}{\sum_{j \in J} \exp(\text{sim}(\text{text}_i, \text{rdf}_j))} \right)$$

$\text{sim}(\text{text}_i, \text{rdf}_j) = \cos(\text{embed}(\text{text}_i), \text{embed}(\text{rdf}_j))$

with  $I$  the set of training instances in the batch. Intuitively, this trains the encoder to learn representations that map text items closer to their RDF anchor than to other RDF graphs in the dataset.

In all our experiments, we start from all-mpnet-base-v2, a pre-trained sentence-MPNet (Song et al., 2020) model, in order to leverage its strong pre-trained text representations.

#### 3.2 Training Datasets

For training, we need  $(g, t)$  pairs where  $g$  is a Wikidata RDF graph and  $t$  is a text in English whose content is similar to  $g$ . We compare three datasets, all created using distant supervision.

**TeKGen.** (Agarwal et al., 2021) use heuristics to align triples from Wikidata to Wikipedia sentences. The TEKGEN dataset covers 1,041 Wikidata properties and consists of about 6M (graph, text) pairs where each text is a sentence.

**KELM.** The KELM corpus has 15M (graph, text) pairs where graphs are created based on relation co-occurrence counts i.e. frequency of alignment of two properties to the same sentence in the training data (Agarwal et al., 2021). Texts are then generated from these graphs using a T5 model fine-tuned on TEKGEN.

**TREx.** (Elsahar et al., 2018) use word- and sentence-tokenization, coreference resolution, a date-time and a predicate linker, plus various RDF-text alignment methods to create TREX, a dataset aligning 11 million Wikidata triples with 6 million Wikipedia sentences.

#### 3.3 Test Datasets

We use two datasets for evaluation: WEBNLG (Gardent et al., 2017) and WIKICHUNKS, which we create in this work. Appendix A shows some statistics for all datasets.

**WebNLG** is a dataset of pairs where the texts were crowdsourced to match the input graph. In WEBNLG the RDF graph is from the DBpedia KB, whereas our models were trained on the Wikidata KB format. To assess the ability of our retrieval model to generalize to different KBs, we evaluate our model both on WEBNLG-DB, the original DBpedia-based dataset, and WEBNLG-WD where the DBpedia graphs have been mapped to Wikidata (Han et al., 2022).

**WikiChunks** consists of 7.3M graph-text pairs where the text is a 100-word *passage* from a Wikipedia dump and the graphs are matching Wikidata graphs. We create matching graphs by aligning all Wikidata  $(s, p, o)$  triples with a Wikipedia passage such that the subject  $s$  of that triple matches




Figure 1: **Retrieval Accuracy** for a variety of training datasets and objectives. Our models outperform the base model (leftmost grey bar) by a large margin. Hard negatives help across the board. Training on an equal mix of datasets yields consistently high performance on aligned (WEBNLG) and noisy (WIKICHUNKS) data.

the entity described by the Wikipedia page from which the passage was extracted and the object  $o$ , or one of its aliases, is mentioned in that passage. Retrieving on this dataset imitates the conditions in which retrieval on Wikipedia is usually executed (Karpukhin et al., 2020b; Lewis et al., 2020). This is a challenging task as, contrary to WEBNLG, WIKICHUNKS matches are not aligned: the Wikidata graph information is strictly included in the passage, which may contain much more. Several passages may also contain very similar information. We use a subset of 30000 pairs, the same size as WEBNLG, to make results comparable.

We evaluate our representations using a retrieval reformulation of the data-to-text NLG task: Given the embedding of a graph, how well can we identify the most similar text in the corpus? As our evaluation sets have 1-to-1 mappings between sources (the graphs) and targets (the texts), the retrieval performance in the opposite direction does not vary by more than 2%. We consider top-result accuracy.

## 4 Results

### 4.1 General Results

We use all-mpnet-base-v2, the state-of-the-art dense sentence embedding model that our models are training from, as a baseline. all-mpnet-base-v2 can estimate semantic similarity, as our models do, but was only trained on text. It can still process the linearized RDF data, however, as it is in the form of natural text. **The baseline is reasonable, but training yields strong improvements** with a top accuracy of 80% for all settings against 38% for the base model (Figure 1)

and 0.003% for random-chance performance.

### 4.2 Generalization to other KB formats

Encoding the RDF data as natural language allows for flexibility in the RDF format, as opposed to earlier graph approaches that encode relations and entities as integers. After fine-tuning on Wikidata graphs, which include relations like `place served by transport hub`, we might be able to generalize to DBpedia, which would use `cityServed` instead, as the base pre-trained model knows all these words. Indeed, we find that **retrieval performance is similar on WEBNLG-WD and WEBNLG-DB**.

### 4.3 Batch Size and Negatives

We experiment with adding artificial hard negatives to the batch, and with different batch sizes. Confounders are constructed from the correct graph by corrupting a triple inside that graph, replacing a subject, object or predicate at random with another subject, object or predicate in the dataset. This form of data augmentation is made possible by the formalized nature of RDF graphs: it would be much harder to create confounders on the text side.

**Hard vs. In-batch negatives** Figure 1 shows retrieval accuracy when using only in-batch vs. using in-batch and hard negatives. We see that hard negatives mostly help when retrieving parallel data (WEBNLG) i.e. when small graph-text mismatches strongly impact accuracy. We also see that hard negatives have the strongest impact on the model trained on TEKGEN, which is also the one with the lowest retrieval accuracy. This suggests that hard negatives are most helpful when the training data



is noisier than the evaluation data.

**Batch size.** As previous work has found that larger batch sizes improve contrastive training (Qu et al., 2021), we experiment with two batch size set-ups: 192<sup>1</sup> and 2560<sup>2</sup>. We do not find that larger batch sizes consistently improve retrieval accuracy, and keep the smaller ones for practical reasons. Figure 8 in appendix B shows detailed results.

#### 4.4 Training Data Quality

The quality of training data has a strong impact on retrieval accuracy. We see that performance varies with the training data used: on WEBNLG retrieval, KELM yields by far the best results followed successively by TREX and TEKGEM. On WIKICHUNKS, which is more loosely aligned, TREX is the best dataset and KELM is slightly behind. We create an equal-mixture dataset by concatenating subsets of equal sizes of each dataset<sup>3</sup>. As the rightmost column in Figure 1 shows, this allows us to capture the best of both worlds. We dub the model trained on this data with hard negatives `all_datasets_hard_negatives`.

The similarity distributions according to `all_datasets_hard_negatives` is shown in Figure 2, which matches those results: KELM is much better aligned. This is in line with intuition as KELM text is generated from the input graphs while TREX and TEKGEM are created using distant supervision. We attempted to bootstrap dataset quality by re-training models on the 50% of the data identified as highest-similarity. We find that this does not increase performance and can even decrease it, probably due to loss of diversity.

#### 4.5 Training Data Quantity

As shown in Figure 3, performance plateaus early in training. The advantage of KELM or the concatenated dataset is not due to their larger size.

### 5 Building a Referenceless Metric for Data-to-text Generation

Commonly-used metrics for Natural Language Generation require references to compare the output against, which must be produced by human annotators. Can we leverage our joint embeddings to compare the output text to the input RDF directly, reducing the necessary resources?

<sup>1</sup>The maximum we could fit on an 8-A100 cloud instance.

<sup>2</sup>The maximum we could fit on a larger cluster.

<sup>3</sup>In total, thrice the size of the smallest dataset, TREX.

#### 5.1 Fine-tuning on Human Judgments of Semantic Adequacy

Our retrieval models can be used to provide a similarity metric between text and formal data in the form of the scalar product or cosine distance in embedding space. We can further improve this metric by fine-tuning on human judgments of RDF-text adequacy. In order to show the generalization strength of this approach, we fine-tune our `all_datasets_hard_negatives` model on human-rated WEBNLG-2017 items, and evaluate on human-rated WEBNLG-2020 items, which uses different test data and different criteria for the assessment of semantic adequacy by human judges.

(Shimorina et al., 2018) provides human judgments for the output of 10 NLG systems from WEBNLG challenge 2017. Each model was evaluated on a sample of 223 texts yielding a total of 2230 generated texts annotated with human judgments for the following three criteria.

- **Semantic adequacy:** Does the text correctly represent the meaning in the data?
- **Grammaticality:** Is the text grammatical (no spelling or grammatical errors)?
- **Fluency:** Does the text sound natural?

(Castro Ferreira et al., 2020) provides human judgments for the output of 16 NLG systems from WEBNLG Challenge 2020. Each model was evaluated on a sample of 178 texts yielding a total of 2,848 generated texts annotated with human judgments for the following five criteria.

- **Data Coverage:** Does the text include descriptions of *all* predicates in the input?
- **Relevance:** Does the text describe *only* triples present in the graph?
- **Correctness:** For graph predicates, does the text correctly describe their arguments?
- **Text Structure:** Is the text grammatical, well-structured, written in acceptable English?
- **Fluency:** Does the text progress naturally and form a coherent, easy-to-understand whole?

We train on the 2017 *semantic adequacy* metric. To assess how well our similarity metric reflects human judgments of similarity between an RDF graph




Figure 2: **Pair similarity distributions** according to `all_datasets_hard_negatives` for all datasets.




Figure 3: **Performance throughout training** evaluated by WEBNLG-WD accuracy. Training for longer than the size of the smallest datasets does not change performance meaningfully.

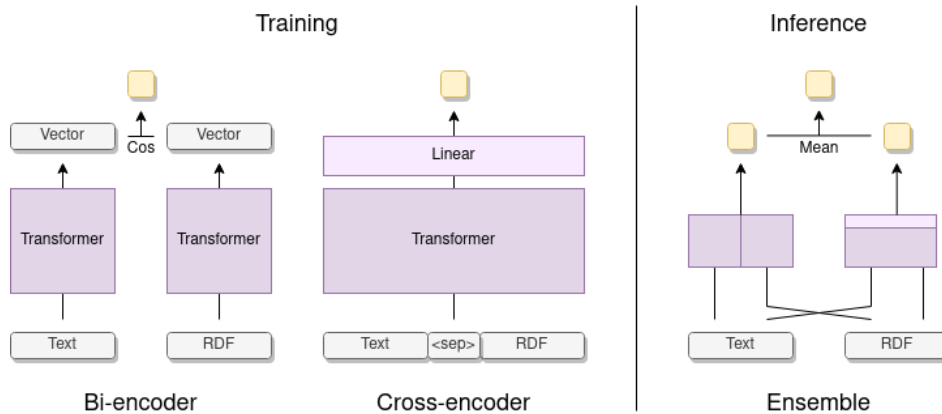


Figure 4: **Fine-tuning setup.** We fine-tune both bi-encoders and cross-encoders on human-rated data. At inference time, we use the mean of a bi-encoder and a cross-encoder as the final metric.

and a Natural Language Text, we compute correlations between our system’s scores and the 2020 human judgments of semantic adequacy, namely *data coverage*, *relevance*, and *correctness*<sup>4</sup>.

## 5.2 Fine-tuning Procedure

**Bi- and Cross-encoder ensembling** We can fine-tune our pre-trained model as a *cross-encoder*, where there is only one instance of the model, which can attend to both items simultaneously and

feed into a linear layer, rather than a *bi-encoder* as previously, where two instances of the model embed the two items separately and the dot product or cosine distance serves as the output. The cross-attention feature allows for higher performance at the cost of making retrieval expensive as all  $n^2$  distances must be computed separately (Humeau et al., 2019). However, bi- and cross-encoders perform well on different data points. The scores they give WEBNLG-2020 candidates have surprisingly low Pearson correlation, 0.66. This makes them good candidates for ensembling, and indeed, taking the mean of the bi- and cross-encoder scores yields

<sup>4</sup>We train on WEBNLG-2017 and evaluate on WEBNLG-2020 as semantic adequacy is a more global criterion encompassing coverage, relevance and correctness while the reverse is not true.




Figure 5: **Difference in similarity between correct and corrupted graph-text pairs.** On the left, `all_datasets_hard_negatives` and `all_datasets_hardinv_negatives` just after pre-training, and on the right, both models after fine-tuning and ensembling on WEBNLG-2017. The system we used as a final metric is the last plot on the right. Models that have seen inverted negatives at pre-training identify correct and corrupted pairs better.

higher correlations with all human judgments. Both architectures and the ensembling method are represented in diagram 4.

**Robustness to inversion** Transformer-based models can sometimes behave as advanced bag-of-word models (Sinha et al., 2021), which would not see a difference if the subject and object are reversed in a triple. In order to examine the robustness, we create an adversarial dataset from all the 1-triple graphs in WEBNLG 2020 with non-symmetrical<sup>5</sup> relationships. In this dataset, for each text, there is a pair with the correct triple and a pair in which the triple’s predicate arguments (subject and object) have been inverted e.g., (*André the Giant, larger than, Samuel Beckett*) vs. (*Samuel Beckett, larger than, André the Giant*). This dataset (WEBNLG-INV) consists of 2793  $(g, t)$ , and  $(g_{inv}, t)$  pairs where  $(g, t)$  is a graph of size one with a non-symmetrical relationship in WEBNLG-WD,  $t$  is the corresponding text and  $g_{inv}$  is the corrupted triple.

We report the difference  $sim(g, t) - sim(g_{inv}, t)$  in the similarity between text and correct graph on the one hand and text and corrupted graph on the other in Figure 5. The higher, the better the model is at recognizing predicate inversion. `all_datasets_hard_negatives`, the retrieval model presented in Section 3.1, does not do well at this task, with 38% of the inverted triplets estimated more similar to the text than the original ones. (After fine-tuning on WEBNLG-2017 judgments, 30%)

<sup>5</sup>Manually defined. The list is in appendix D.

In order to make our models robust to inversion, at pre-training time, we add inverted negatives to the mix of artificial negatives in the batches: confounding graphs where a random triplet has been inverted. The resulting model, `all_datasets_hardinv_negatives` has the same retrieval accuracy but gains inversion detection abilities. This ability is conserved through fine-tuning, as Figure 5 shows: only 14% of triplets are misclassified.

**The final system we choose as a metric** is the ensemble of a bi- and cross-encoder pre-trained on the concatenation of KELM, TEKGEN and TREX with our two types of data augmentation, then fine-tuned on WEBNLG-2017 human judgments. We call it EREDAT, for Ensembled Representations for Evaluation of DATA-to-Text.

### 5.3 Comparison with other Evaluation Metrics

Correlations with human judgments are shown in Figure 6 for a variety of automated evaluation metrics: three metrics that require a reference (BLEU, BERTscore-F1, and BLEURT, the previous state of the art) and two referenceless metrics (DataQuestEval and EREDAT). Our metric is the best correlated with all human judgment categories, even including metrics with references. As shown in 7, this advantage is mostly explainable by EREDAT’s improved robustness to longer, more complex graphs, which tend to degrade correlation with human judgment. Scatter plots of the underlying distributions are given in appendix C.

As human references are rarely available and




Figure 6: **Pearson correlation between automatic metrics and human judgments.** Lighter and higher is better. EREDAT outperforms the other referenceless metric and matches BLEURT, which requires a reference.




Figure 7: **Correlation with human judgment by graph size** for EREDAT, DQE and BLEURT. Our metric is more robust than BLEURT to longer graphs, and generally much more correlated than DQE, the existing referenceless metric.

costly to produce, and EREDAT attains higher correlation with human judgments without relying on them, it is the most practical choice to evaluate data-to-text generation. In this case, it was not fine-tuned to the same kind of data it was applied to, showing it generalizes to new datasets. If one has a specific dataset or task in mind, even better performance could be attained by training on a set of problem-specific human judgments.

## 6 Conclusion

We presented an architecture and pre-training strategy to measure the similarity between RDF graphs and English texts, introducing novel data augmentation strategies made possible by the RDF structure.

Specifically, we introduced a bi-encoder retrieval model trained on unlabeled RDF-text data which achieves high retrieval accuracy on both parallel and real-life, less well aligned datasets. Building from this pre-trained model, we further provided a novel evaluation metric for RDF-to-text generation models which matches state-of-the-art metrics in terms of correlation with human judgments of semantic adequacy without needing costly human-written references. This metric can also be used to filter existing text/RDF datasets.

## 7 Limitations and compute statement

This study focuses on English text. Reproducing the proposed approach for use on other languages would require dedicated datasets of similar scale, along with graph/text alignments. Further, the other languages differ quite a lot from English-centric RDF graphs, potentially reducing the suitability of the proposed framework and requiring more advanced multilingual methods. We release our models with the intended use of representation learning and automated RDF-to-text evaluation. Other uses may not be appropriate.

We trained over 2000 models for a total of approximately 2400 GPU-hours (NVIDIA V100s and A100s) of compute on public infrastructure and Google Compute Engine. Most of them were based on all-mpnet-base-v2, with 109M parameters.



## 8 Acknowledgments

We thank the anonymous reviewers for their feedback. We gratefully acknowledge the support of the French National Research Agency (Gardent; award ANR-20-CHIA-0003, XNLG "Multilingual, Multi-Source Text Generation").

## References

- Oshin Agarwal, Heming Ge, Siamak Shakeri, and Rami Al-Rfou. 2021. [Knowledge graph based synthetic corpus generation for knowledge-enhanced language model pre-training](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3554–3565, Online. Association for Computational Linguistics.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Antoine Bordes, Nicolas Usunier, Alberto Garcia-Durán, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. In *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 2, NIPS'13*, page 2787–2795, Red Hook, NY, USA. Curran Associates Inc.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Thiago Castro Ferreira, Claire Gardent, Nikolai Ilinykh, Chris van der Lee, Simon Mille, Diego Moussallem, and Anastasia Shimorina. 2020. [The 2020 bilingual, bi-directional WebNLG+ shared task: Overview and evaluation results \(WebNLG+ 2020\)](#). In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 55–76, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Jiecao Chen, Liu Yang, Karthik Raman, Michael Bendersky, Jung-Jung Yeh, Yun Zhou, Marc Najork, Danyang Cai, and Ehsan Emadzadeh. 2020. [DiPair: Fast and accurate distillation for trillion-scale text matching and pair modeling](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2925–2937, Online. Association for Computational Linguistics.
- Muhao Chen, Yingtao Tian, Kai-Wei Chang, Steven Skiena, and Carlo Zaniolo. 2018. [Co-training embeddings of knowledge graphs and entity descriptions for cross-lingual entity alignment](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 3998–4004. International Joint Conferences on Artificial Intelligence Organization.
- Tim Dettmers, Pasquale Minervini, Pontus Stenetorp, and Sebastian Riedel. 2018. Convolutional 2d knowledge graph embeddings. In *Thirty-second AAAI conference on artificial intelligence*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ondřej Dušek and Zdeněk Kasner. 2020. [Evaluating semantic accuracy of data-to-text generation with natural language inference](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 131–137, Dublin, Ireland. Association for Computational Linguistics.
- Hady Elsahar, Pavlos Vougiouklis, Arslan Remaci, Christophe Gravier, Jonathon Hare, Frederique Laforest, and Elena Simperl. 2018. [T-REx: A large scale alignment of natural language with knowledge base triples](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. [Creating training corpora for NLG micro-planners](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 179–188, Vancouver, Canada. Association for Computational Linguistics.
- Kelvin Han, Thiago Castro Ferreira, and Claire Gardent. 2022. [Generating questions from Wikidata triples](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 277–290, Marseille, France. European Language Resources Association.
- Xu Han, Zhiyuan Liu, and Maosong Sun. 2016. Joint representation learning of text and knowledge for knowledge graph completion. *ArXiv*, abs/1611.04125.
- Matthew L. Henderson, Rami Al-Rfou, Brian Strope, Yun-Hsuan Sung, László Lukács, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. 2017. [Efficient natural language response suggestion for smart reply](#). *CoRR*, abs/1705.00652.

- Samuel Humeau, Kurt Shuster, Marie-Anne Lachaux, and Jason Weston. 2019. [Real-time inference in multi-sentence tasks with deep pretrained transformers](#). *CoRR*, abs/1905.01969.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020a. [Dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6769–6781, Online. Association for Computational Linguistics.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020b. [Dense passage retrieval for open-domain question answering](#). *CoRR*, abs/2004.04906.
- Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive NLP tasks](#). *CoRR*, abs/2005.11401.
- Tomaš Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#).
- Eric Miller. 1998. An introduction to the resource description framework. *D-lib Magazine*.
- Maximilian Nickel, Lorenzo Rosasco, and Tomaso Poggio. 2016. Holographic embeddings of knowledge graphs. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, AAAI'16*, page 1955–1961. AAAI Press.
- Maximilian Nickel, Volker Tresp, and Hans-Peter Kriegel. 2011. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11*, page 809–816, Madison, WI, USA. Omnipress.
- Vardaan Pahuja, Yu Gu, Wenhua Chen, Mehdi Bahrami, Lei Liu, Wei-Peng Chen, and Yu Su. 2021. [A systematic investigation of KB-text embedding alignment at scale](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1764–1774, Online. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Pouya Pezeshkpour, Liyan Chen, and Sameer Singh. 2018. [Embedding multimodal relational data for knowledge base completion](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3208–3218, Brussels, Belgium. Association for Computational Linguistics.
- Yingqi Qu, Yuchen Ding, Jing Liu, Kai Liu, Ruiyang Ren, Wayne Xin Zhao, Daxiang Dong, Hua Wu, and Haifeng Wang. 2021. [RocketQA: An optimized training approach to dense passage retrieval for open-domain question answering](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5835–5847, Online. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *ICML*.
- Clement Rebuffel, Thomas Scialom, Laure Soulier, Benjamin Piwowarski, Sylvain Lamprier, Jacopo Staliano, Geoffrey Scuttheeten, and Patrick Gallinari. 2021. [Data-QuestEval: A referenceless metric for data-to-text semantic evaluation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8029–8036, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Michael Schlichtkrull, Thomas Kipf, Peter Bloem, Rianne Berg, Ivan Titov, and Max Welling. 2018. [Modeling Relational Data with Graph Convolutional Networks](#), pages 593–607.
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. [Facenet: A unified embedding for face recognition and clustering](#). *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020a. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Thibault Sellam, Amy Pu, Hyung Won Chung, Sebastian Gehrmann, Qijun Tan, Markus Freitag, Dipanjan Das, and Ankur Parikh. 2020b. [Learning to evaluate translation beyond English: BLEURT submissions](#)

- to the WMT metrics 2020 shared task. In *Proceedings of the Fifth Conference on Machine Translation*, pages 921–927, Online. Association for Computational Linguistics.
- Anastasia Shimorina, Claire Gardent, Shashi Narayan, and Laura Perez-Beltrachini. 2018. [WebNLG Challenge: Human Evaluation Results](#). Technical report, Loria & Inria Grand Est.
- Koustuv Sinha, Robin Jia, Dieuwke Hupkes, Joelle Pineau, Adina Williams, and Douwe Kiela. 2021. [Masked language modeling and the distributional hypothesis: Order word matters pre-training for little](#). *CoRR*, abs/2104.06644.
- Matthew Snover, Bonnie Dorr, Rich Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA. Association for Machine Translation in the Americas.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. [Mpnet: Masked and permuted pre-training for language understanding](#). *CoRR*, abs/2004.09297.
- Kristina Toutanova, Danqi Chen, Patrick Pantel, Hoi-fung Poon, Pallavi Choudhury, and Michael Gamon. 2015. [Representing text for joint embedding of text and knowledge bases](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509, Lisbon, Portugal. Association for Computational Linguistics.
- Théo Trouillon, Johannes Welbl, Sebastian Riedel, Éric Gaussier, and Guillaume Bouchard. 2016. [Complex embeddings for simple link prediction](#). In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, page 2071–2080. JMLR.org.
- Zhigang Wang and Juanzi Li. 2016. [Text-enhanced representation learning for knowledge graph](#). In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI'16*, page 1293–1299. AAAI Press.
- Sam Wiseman, Stuart Shieber, and Alexander Rush. 2017. [Challenges in data-to-document generation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2253–2263, Copenhagen, Denmark. Association for Computational Linguistics.
- Jiawei Wu, Ruobing Xie, Zhiyuan Liu, and Maosong Sun. 2016. [Knowledge representation via joint learning of sequential text and knowledge graphs](#). *arXiv preprint arXiv:1609.07075*.
- Kun Xu, Liwei Wang, Mo Yu, Yansong Feng, Yan Song, Zhiguo Wang, and Dong Yu. 2019. [Cross-lingual knowledge graph alignment via graph matching neural network](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3156–3161, Florence, Italy. Association for Computational Linguistics.
- Ikuya Yamada, Hiroyuki Shindo, Hideaki Takeda, and Yoshiyasu Takefuji. 2016. [Joint learning of the embedding of words and entities for named entity disambiguation](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 250–259, Berlin, Germany. Association for Computational Linguistics.
- Bishan Yang, Wen-tau Yih, Xiaodong He, Jianfeng Gao, and li Deng. 2014. [Embedding entities and relations for learning and inference in knowledge bases](#).
- Haiyang Yu, Ningyu Zhang, Shumin Deng, Hongbin Ye, Wei Zhang, and Huajun Chen. 2020. [Bridging text and knowledge with multi-prototype embedding for few-shot relational triple extraction](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6399–6410, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tianyi Zhang\*, Varsha Kishore\*, Felix Wu\*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.

## A Dataset statistics

	# (t,g)	# P	# E
TEKGEN	6,310,061	1041	3,939,696
TREX	6,000,336	675	3,188,309
KELM	15,616,551	261405	5,073,603
WEBNLG-DB	13,212	372	3210
WEBNLG-WD	10,384	188	2783
WIKICHUNKS	30,000	468	20,318

Table 1: **Training and test data for retrieval.** # (t,g): Number of graph-text pairs, # T: Number of texts, # G: Number of graphs, # P: Number of distinct properties, # E: Number of distinct entities.

## B Impact of Batch Size

## C Scatter Plot Comparison of BLEURT and EREDAT

## D Symmetrical Relationships in WebNLG

We manually inspected all relationships in WEBNLG and deemed the following to be symmetrical in nature:

"taxon synonym", "partner in business or sport", "opposite of", "partially coincident with", "physically interacts with", "partner", "relative", "related category", "connects with", "twinned administrative body", "different from", "said to be the same as", "sibling", "adjacent station", "shares border with"





Figure 8: **Small vs. Large Batch Size.** Large batch sizes help a little on data with lower alignment quality (WIKICHUNKS). Overall, the improvement is inconsistent.



Automatic metrics vs. function of human judgments in WebNLG 2020

Figure 9: **Human judgment and automated evaluation values for every point in WEBNLG 2020.**