

Denoising Pre-Training and Data Augmentation Strategies for Enhanced RDF Verbalization with Transformers

Sebastien Montella*

Orange Labs, Lannion

sebastien.montella@orange.com

Betty Fabre*

Orange Labs, Lannion

betty.fabre@orange.com

Tanguy Urvoy

Orange Labs, Lannion

tanguy.urvoy@orange.com

Johannes Heinecke

Orange Labs, Lannion

johannes.heinecke@orange.com

Lina Rojas-Barahona

Orange Labs, Lannion

linamaria.rojasbarahona@orange.com

Abstract

The task of verbalization of RDF triples has known a growth in popularity due to the rising ubiquity of Knowledge Bases (KBs). The formalism of RDF triples is a simple and efficient way to store facts at a large scale. However, its abstract representation makes it difficult for humans to interpret. For this purpose, the WebNLG challenge aims at promoting automated RDF-to-text generation. We propose to leverage pre-trainings from augmented data with the Transformer model using a data augmentation strategy. Our experiment results show a minimum relative increases of 3.73%, 126.05% and 88.16% in BLEU score for seen categories, unseen entities and unseen categories respectively over the standard training.

1 Introduction

The purpose of Knowledge Bases is to store unstructured data into a well-defined and highly structured representation. A commonly used formalism is the Resource Description Framework (RDF) where knowledge takes the form of triplets (subject; predicate; object). For instance if we take the triple $\langle \text{Alan_Bean}; \text{birthDate}; 1932 \rangle$, the subject is *Alan_Bean*, the predicate/relation/property is *birthDate* and the object is the date *1932*. This triple-based representation allows one to describe complex relations between entities. It is hence frequently used in natural language processing (NLP) applications such as question answering, summarization, recommendation or search engines.

Although the RDF format is an efficient method to store knowledge, triples are hard to interpret by humans. The automatic translation from sets of triples into human language is however a difficult task. Thanks to the WebNLG challenge (Gardent et al., 2017b), this RDF-to-text task attracts now many researchers. This challenge aims at generating a faithful and descriptive textual content from RDF inputs. Given the triples $\langle \text{Alan_Bean}; \text{birthPlace}; \text{Wheeler, Texas} \rangle$ and $\langle \text{Alan_Bean}; \text{birthDate}; 1932-03-15 \rangle$, one possible verbalization is “*Alan Bean was born in Wheeler, Texas on Mar 15, 1932.*”. Verbalization of structured data can provide an unambiguous and simple translation from abstract representations.

The WebNLG dataset is frequently used as a benchmark for Information Extraction (IE), Natural Language Understanding (NLU) or Natural Language Generation (NLG) tasks. It offers different lexicalization for the same set of triples. Despite limited in size, the high-quality of both its textual content and data-to-text mapping significantly stimulated both RDF-to-text, *i.e.* NLG, and text-to-RDF, *i.e.* IE and NLU, research directions. For instance, Iso et al. (2020) and Gao et al. (2020) have obtained state-of-the-art performances on the WebNLG RDF-to-text task. On the other hand, Guo et al. (2020) considered both NLG and IE objectives within the same learning framework via cycle training. Furthermore, the exploration of a new evaluation metric is also investigated on this WebNLG dataset, confirming its attractiveness for the research community (Dhingra et al., 2019; Sel-

*Both authors contributed equally.

lam et al., 2020).

Recently, the field of NLP has been significantly impacted by large pre-trained language models and novel learning architectures that improved state-of-the-art results on several benchmarks (Devlin et al., 2019; Zhang et al., 2019; Yang et al., 2019; Radford et al., 2019; Lewis et al., 2020; Clark et al., 2020). The Transformer model (Vaswani et al., 2017) is in most cases at the core of these improvements. The self-attention mechanism of Transformer enables a better understanding of underlying dependencies in sentences.

However, the aforementioned techniques have not yet been tested on the RDF-to-text task. We therefore propose to combine the Transformer model with large pre-trainings to generate more accurate verbalization of the RDF triples. We chose to make use of external corpora during pre-training to reach a better generalization. We build two additional datasets for denoising pre-training and data augmentation purposes. This is motivated by the fact that participants’ systems at previous WebNLG challenge suffered from high drops in performance when evaluated on new entities and predicates not encountered during the training phase (Gardent et al., 2017b). Thus, by incorporating external corpora, we expect our system to outperform systems solely trained on the WebNLG dataset. We evaluate our approach using BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), chrF++ (Popović, 2015) and BLEURT (Sellam et al., 2020). As mentioned in (Post, 2018), BLEU is a parameterized metric that may exhibit wild variations. The use of multiple metrics gives better outlooks on the general performance.

The reminder of this paper is organized as follows. Section 2 gives an overview of existing methods to translate RDFs into their corresponding textual content. After introducing the task and the WebNLG datasets in Section 3, we introduce our approaches for the challenge in Section 4. The different training phases are presented in Section 5. We show our results in Section 6. Finally, we conclude in Section 7.

2 Related Work

Following the traditional NLG pipeline (Reiter and Dale, 1997), early works on factual data verbalization mainly focused on the creation of rules or templates to produce the textual output. Sun and Mellish (2006, 2007) have shown that most of the

useful information for the generation is brought by the rich linguistic information of the RDF properties, *i.e.* predicates. They have noticed that properties can be classified into 6 categories based on their pattern. For instance, some predicates may be a concatenation of a verb and a noun like *hasEmail*, or starting with a verb and finishing with an adjective like *hasTimeOpen*. They established handcrafted rules to construct linguistic forms for each category. This technique allows a domain independent verbalization since relying only on the pattern of predicates. On the contrary, Cimiano et al. (2013) used domain-dependent ontology lexicon to proceed to a fine-grained and specific verbalization of mentioned concepts. Although some methods tried to automatically learn those templates (Duma and Klein, 2013), a major disadvantage of such approaches is the need of handcrafted rules and the poor generalization power.

The main downside of KB verbalization is the lack of paired data. In spite of the impressive success of Deep Learning (DL) methods, few attempts could be made toward RDF-to-text task due to data scarcity. To address this issue, Gardent et al. (2017a) came up with a new dataset which includes RDFs extracted from DBPedia¹ with their corresponding translations in natural language. To further stimulate the development of KB verbalizer, Gardent et al. (2017b) set up the WebNLG challenge 2017, welcoming new techniques for the RDF-to-text task. Among the proposed strategies, DL-based models outperformed rules-based systems (Gardent et al., 2017b). Neural approaches relied almost on the attention-based encoder-decoder architecture (Sutskever et al., 2014; Bahdanau et al., 2015). The encoder-decoder framework showed promising results, with some limitations on unseen domains.

Following WebNLG 2017, multiple works were conducted to cope with the witnessed drawbacks of participants’ system. Zhu et al. (2019) argued that Kullback-Leibler (KL) divergence should be avoided. Minimizing the KL divergence contributes to more diversity in generated samples, at the cost of quality. As a matter of fact, diversity is unnecessary for KBs verbalization since it is constrained on specific entities and relations. Therefore, they proposed to optimize the inverse KL divergence to generate higher-quality textual description. Another flaw appears in the linearization

¹<https://wiki.dbpedia.org/>

of input triples. Such transformation gets rid of the intrinsic structure of the RDFs graph which may lead to inconsistency or relational mismatch in predicted sentences. To address this question and upon the rising popularity of Graph Neural Networks (GNNs), graph-based solutions were presented. [Trisedya et al. \(2018\)](#) introduced a GTR-LSTM triple encoder. The idea is to traverse the graph in a specific manner such that vertices are better ordered to leverage the *intra*- and *inter*-triples relationships.

More recently, Graph Neural Networks (GNNs) have raised a lot of attention to capture dependencies between nodes of a graph ([Zhou et al., 2018](#)). The message passing scheme offers the ability to better model relations nodes and their neighborhoods. Well-suited for the generation from structured data, [Marcheggiani and Perez-Beltrachini \(2018\)](#) made use of a Graph Convolutional Network (GCN) to explicitly encode the input structure of the triples instead of a straightforward linearization. More precisely, they used an encoder-decoder architecture composed of a GCN encoder and a LSTM decoder. They further applied a copy mechanism from ([See et al., 2017](#)) to copy words from the input RDF triples. The resulting performances confirmed the benefit of the graph-based encoder. Following this conclusion, several papers proposed graph-based encoder as a solution ([Gao et al., 2020](#); [Zhao et al., 2020](#); [Moussallem et al., 2020](#)).

Meanwhile, in contrast to the conditional text generation paradigm, [Iso et al. \(2020\)](#) considered the KB verbalization as a text editing task with their FACTEDITOR model. Text editing consists in applying a sequence of transformations on an input text to produce a particular output. FACTEDITOR, a simple BiLSTM model, creates a verbalization candidate, so-called *draft text*, from which the edition process starts. Three actions can be made on a candidate word: *keeping* it, *dropping* it or *generating* a new one.

Recent studies also explore the fusion of both NLU and NLG into one unified model. The RDF-to-Text and Text-to-RDF can be seen as the inverse task from one another. Inspired by the advancement from computer vision ([Zhu et al., 2017](#)), [Tseng et al. \(2020\)](#) and [Guo et al. \(2020\)](#) provide a cycle training framework to learn both tasks simultaneously.

Nevertheless, none of the above techniques use large pre-training neither take advantage of the Transformer model. We propose in this paper to in-

vestigate the effect of massive pre-training coupled with the vanilla Transformer for the RDF-to-Text task.

3 The WebNLG Challenge

In this section we present the WebNLG task and datasets.

3.1 Task Definition

We call *RDF verbalization* and *RDF lexicalization* the mapping of the RDFs triples to their descriptive text. It is a variant of the machine translation task where the input is a list of triples instead of a raw text. Let T be a set of RDFs triples of size N , with t_i the i^{th} triple, $i \in [1 \dots N]$. We write each triple t_i as $\langle s_i; p_i; o_i \rangle$ with s_i , p_i and o_i the subject, predicate and object of the i^{th} triple respectively. The k^{th} verbalization of T is denoted v^k . The RDF-to-text task consists in mapping the set T to one of its possible lexicalization v_k . The generated verbalization should be grammatically and semantically correct. Regarding the knowledge, it should also contain the same level of information as the input triples.

3.2 Datasets

Initially, [Gardent et al. \(2017a,b\)](#) fashioned a corpus filled with *triple-sets*, *i.e.* sets of RDF triples, and their possible lexicalizations, both in English. Each tripletset is related to a particular DBpedia category (*e.g.* *Astronaut*) and can contain up to 7 triples. For this new version of the WebNLG 2020, some changes were carried out to encourage participation to the competition. First, the unseen testing categories of the original WebNLG dataset are now part of the training set. The new extra category *Company* is also included reaching now 16 domains for training. Moreover, to go beyond monolingualism, the second WebNLG challenge propose a bilingual corpus by inserting Russian translations for some RDF triples beside English ones. In our work, we however only focus on English verbalization. The support of multiple languages is part of our for future work.

The training set contains 13,211 unique triple-sets for which different possible lexicalization can be given. The more triples in a tripletset, the higher the number of possible lexicalization. Totally, we can list out 35,426 RDFs-Text pairs. Note that a corresponding verbalization of a set of RDFs is not necessary a single sentence. Multiple sen-

tences may encompass the RDFs information and co-references may therefore appear. The official WebNLG testing set is divided into three parts for fine-grained evaluation: seen categories, unseen entities, and unseen categories subsets. We give the evaluation performance on those subsets in Section 6.3.

4 Proposed Approaches

We propose to use a Transformer model with data augmentation.

4.1 The Transformer Model

The Transformer model introduced by Vaswani et al. (2017) is an encoder-decoder neural network architecture that relies on the self-attention mechanism. In an encoder-decoder architecture, the encoder maps the input sequence to a sequence of continuous representations from which the decoder generates an output sequence (Sutskever et al., 2014). The Transformer model makes use of stacked layers of multi-head self-attention in both the encoder and decoder.

Self-attention allows longer dependencies to be handled than the seq2seq model that relies on recurrent LSTM or GRU cells (Sutskever et al., 2014; Bahdanau et al., 2015). The self-attention mechanism seems well-suited to handle subject, predicate, and object dependencies within the RDF itself but also dependencies between triples.

As highlighted in the original paper (Vaswani et al., 2017), the individual attention heads from the multi-head attention layers learn to perform individually different tasks. In particular, they appear to learn the syntactic and semantic structure of the sentences independently. Mareček and Rosa (2019) has shown that the self-attention layers embed implicitly syntactic parse trees. Similarly, we hope that our model benefits from our pre-training stage by incorporating such syntactic information, while learning intra and inter-triples dependencies in the fine-tuning stage.

4.2 Data Augmentation

A major shortcoming of the WebNLG dataset remains its size. The lack of diversity impinges on the generalization of unseen categories. We report only 747 and 385 distinct subject entities and predicates, respectively. Thus, our intuition is to enlarge our RDF-to-text mappings to handle multiple relations between different type of entities. To do

so, we proceed in two phases. First, we collect a large set of sentences from which, in our second stage, triples will be extracted. We hereinafter give details for each stage.

4.2.1 Sentence Collection

We make use of Wikipedia that offers a massive and free amount of data to be leveraged. The WebNLG dataset was constructed from DBpedia database which itself relies on Wikipedia. The substantial advantage of including new Wikipedia content lies in that new named entities can then be encountered from much more heterogeneous domains. This allows for a better generalization and knowledge integration. As observed in language models (Devlin et al., 2019; Radford et al., 2019; Zhang et al., 2019), knowledge is somehow assimilated by the model throughout training.

We gathered 13,614 Wikipedia pages containing about 103 million sentences totally². To alleviate the training process, we filter sentences that may impede generation quality. We discard a sentence (1) if it does not start with a upper case letter and does not end with a period, (2) if its length is longer than 50 or (3) if it contains special characters. After this first selection, 57 million (unique) sentences remain. Henceforth, we coin this set of sentences WS1.

4.2.2 Triple Extraction

Triples extraction is challenging and non-trivial, demanding a sharp understanding of linguistic structure. For simplicity (and time-saving), we take advantage of the Stanford Open Information Extraction (Stanford OpenIE) tool³ to extract triples from WS1 sentences. Stanford OpenIE package is schema-free. That means no preliminary definition of the possible predicates is required, as opposed to usual RDFs extractor. The raw text linking two entities will be retrieved as the predicate. Unfortunately, the returned triples may be incomplete, false or alike. For example, for the input sentence "Barack Obama was born in Hawaii.", the returned triples by Stanford OpenIE are $\langle \text{Barack Obama}; \text{was}; \text{born} \rangle$ and $\langle \text{Barack Obama}; \text{was born in}; \text{Hawaii} \rangle$. The first output triple is spurious besides expressing somehow the same insight than the second triple. As a result, a filtering step is crucial to reduce false

²We used the 20200401 Wikipedia dump ~ 18GB

³<https://nlp.stanford.edu/software/openie.html>

triples. We fashion simple rules to limit redundancy by comparing each retrieved triple with each other. Let t_i and t_j be two retrieved triples, with $i \neq j$. We linearize each triple by concatenating words from the subject, predicate and object. The linearized version of t_i and t_j are named t_i^l and t_j^l , respectively. To detect if t_i^l and t_j^l are equivalent, we first verify if either $t_i^l \in t_j^l$ or $t_j^l \in t_i^l$. We further compute the *Levenshtein distance* (Levenshtein, 1965) to consider minor variations between triples. Two triples with an edit distance smaller or equal to 2 are also believed to be similar. If triples similarity is confirmed, the longest triple is kept because sharing the more information with the input sentence. The higher the lexical coverage, the better. Indeed, the risk is to miss essential information from the sentence. If input triples are considerably incomplete, the model will suffer from hallucination, *i.e.* generating content not present in the given input.

However, when t_i^l and t_j^l are too different but semantically identical, more sophisticated conditions should be exploited. For this purpose, we make use of BLEU (Papineni et al., 2002). Extracted triples are derived from the same sentence. Therefore, a *n-gram* based metric such as BLEU is a good choice to check if two triples are alike. We assume triples to be analogous if their BLEU score is greater than 50. In such case, we keep the triple maximizing its BLEU score with the reference sentence for better coverage, as explained previously. The final collection of remaining triples-sentences pairs is called ST1.

Nevertheless, the detection of erroneous information conveyed by the triples remain hard to operate. We are aware that mistaken triples jeopardize directly the performance of our model. That is why we leverage these augmented data in a pre-training step, as detailed in Section 5.1.

5 Training

We suggest in our work to thoroughly combine a pre-training stage before training our Transformer model on the WebNLG dataset. We describe our learning design in the following subsections.

5.1 Pre-Training Objectives

From word embeddings to recently massive language models, NLP has undoubtedly benefited from large pre-training. Many downstream tasks could leverage those models by an effortless fine-

tuning. Recent pre-trained models like BART (Peters et al., 2019) harness unlabeled data to boost model performance in a self-supervised setting. The denoising autoencoder objective from BART has shown a significant performance gain. A denoising objective aims to reconstruct a corrupted input. When fine-tuned on a text generation task, BART has confirmed to be highly effective. We therefore adopt a similar approach using our voluminous Wikipedia sentences, *i.e.* WS1 dataset. Peters et al. (2019) use an arbitrary noising function to permute, delete, and mask words in input. Our transformations differ from BART in that we want our corrupted input to contain factual information, similarly to what our model will be exposed to when fine-tuned on our RDF-to-text task. With this in mind, we choose to keep words with specific Part-Of-Speech (POS) tags like nouns, verbs, adjectives and adverbs. Words with other tags are removed. Modal verbs (*e.g.* *should*) and passive forms are ignored as well. As an example, let the Wikipedia sentence “*In 1860 few of the streets north of 42nd had been graded.*”. After our noising transformations, we obtain the corrupted input “*1860 few streets north 42nd graded.*”. We observe that remaining words stand for the *semantic mass* of the sentence from which we need to lexicalize and connect entailed concepts and entities properly.

At the difference of RDFs triples, the corrupted sentences have no inherent structure. In order to avoid too much divergence between pre-training and fine-tuning inputs’ representation, we follow the denoising pre-training with another pre-training process on a RDF-to-text task. To do so, We use our constructed ST1 dataset which contains triple-based inputs, as detailed in Section 4.2. We expect that these two pre-training steps will improve our ability to generate faithful and consistent verbalization for the WebNLG Challenge.

5.2 Fine-tuning Settings

We coerce our model to use a WebNLG-based vocabulary during pre-training so that the same model can be straightly fine-tuned without any vocabulary discrepancy. Technically, we load the last checkpoint from pre-trained model and directly start to fine-tune it by resetting the optimizer and setting the new data loader to the WebNLG training set. We also tried without resetting the optimizer and obtained similar performances.

In addition, we apply a Curriculum Learning

(CL) approach for faster convergence and better local minima through the optimization process. CL was proposed by Bengio et al. (2009), motivated by the learning process of humans and animals. The core idea is to gradually increase the complexity of samples seen during training instead of a random selection. In our case, we define complexity as the number of RDF triples in the input, *i.e.* N . The higher the number of triples, the higher the complexity. We sort the WebNLG pairs such that easier examples come first and then harder examples afterwards. Our curriculum approach differs in that both easier and harder examples are seen within the same epoch, but in a gradually order. In (Bengio et al., 2009), more complex samples are progressively added during training. To witness the effect of increasing complexity while training, we fine-tune our pre-trained models with and without curriculum learning.

6 Experiments

We evaluate our models on the WebNLG Dataset test set. We compare 8 Transformers models corresponding to 8 different training settings.

6.1 Training and decoding settings

In the following, we present the preprocessing, the training phase and the decoding pipeline.

Preprocessing For preprocessing, we used the Moses tokenizer⁴ and subword segmentation following Sennrich et al. (2016) with the subword-nmt library⁵. We left both the input triples and output text true-cased.

In the context of the WebNLG challenge, we considered a transduction strategy (Gammerman et al., 1998). We restricted the vocabulary to the WebNLG dataset (training, validation and testing set). This learning scheme aims at performing well on a specific set, and not necessary to generalize.

The training process remained the same for all our experiments. Only the data loaded in the batches during training were changed.

To deal with the RDF triples format, we added to the vocabulary four special tokens, namely $\langle object \rangle$, $\langle subject \rangle$, $\langle predicate \rangle$ and $\langle eot \rangle$ (end of triple), that we used as separators within and between triples.

⁴Available on github <https://github.com/moses-smt/mosesdecoder>

⁵Available on github <https://github.com/rsennrich/subword-nmt>

In the case of multiple triples, we built the Transformer input sequence by concatenating triples one after the other. We used this input format on both ST1 and WebNLG.

For the denoising pre-training, we used the WS1 dataset. The transformed sentences (see Section 5.1) were directly fed to the Transformer. The model thus has to reconstruct the incomplete sentences.

Examples of the preprocessed samples for the different training settings are reported in Table 1.

Training We used the Transformer implementation from FAIRSEQ library (Ott et al., 2019)⁶ with the *transformer_base* hyper-parameters set from (Vaswani et al., 2017). We optimized the weights of our neural networks using an ADAM optimizer and a label smoothed cross entropy loss.

We made 10 epochs of pre-training and stopped fine-tuning when the performance with the BLEU score on the validation set did not improved after 30 epochs. At the end, we kept the model that achieved the best BLEU on the validation set.

To study the impact of curriculum learning, we launched fine-tuning with and without it. In the former case, we prevented the shuffling of the batches for 30 epochs. The data being sorted by number of triples, the model had to deal first with simple samples, and then with more complex ones as the learning did progress. In the results, we report this setting as CL (for Curriculum Learning).

Decoding For decoding, we did a beam search with a beam of size 5. We merged the subwords back into words and detokenized.

6.2 Evaluation setting

For evaluation, we used the official WebNLG evaluation script⁷. The metrics we used to compare our models are BLEU, METEOR, chrF++ and BLEURT. BLEURT metric has been recently proposed and proven to be well-correlated with human judgments. Relying on BERT’s contextual embeddings, BLEURT offers semantically robust feedback. The *n-gram*-based evaluation techniques such as BLEU, METEOR or chrF++ are additional metrics to judge the generation quality. When used together, they give good assessment of the generation quality of our system.

⁶We used the compiled version 0.9.0 from <https://github.com/pytorch/fairseq>

⁷Available on github <https://github.com/WebNLG/GenerationEval>

Task	Input sequence	Output sequence
WebNLG	$\langle subject \rangle$ Italy $\langle predicate \rangle$ capital $\langle object \rangle$ Rome $\langle eot \rangle$	Rome is the capital of Italy .
	$\langle subject \rangle$ Bionico $\langle predicate \rangle$ course $\langle object \rangle$ Dessert $\langle eot \rangle$ $\langle subject \rangle$ Bionico $\langle predicate \rangle$ ingredient $\langle object \rangle$ Raisin $\langle eot \rangle$	Bionico is a dessert which contains raisins .
WS1	He died he@@ art f@@ ail@@ ure h@@ ospital October 5 2014	He died of he@@ art f@@ ail@@ ure at the h@@ ospital on October 5 , 2014 .
ST1	$\langle subject \rangle$ He $\langle predicate \rangle$ retired $\langle object \rangle$ 199@@ 0 $\langle eot \rangle$	He retired in 199@@ 0 .

Table 1: **Preprocessed samples** fed to the Transformer model during pre-training or fine-tuning. The triples are concatenated into one sequence. We used Moses tokenizer and a subwords encoding, the subwords are divided by @@ . Special tokens are used to keep the RDFs triplet format : $\langle subject \rangle$, $\langle predicate \rangle$, $\langle object \rangle$ and $\langle eot \rangle$. In the case of the denoising task (**WS1**) we do not have these special tokens. We left the data truncated and no delexicalization is applied.

Unlike in 2017, this time the generated sentences must be detokenized and true-cased. Our models are therefore not directly comparable with the models of the previous version of the challenge.

Also, as the test base is not available, we do not know how many references our generated sentences are compared to.

6.3 Results

To gauge the influence of our different learning approaches, we conducted an ablation study. We defined our baseline as the Transformer solely trained on WebNLG dataset, without curriculum. We want to assess the influence of pre-training, data augmentation and curriculum learning compared to the baseline. To provide a fair and detailed analysis, we evaluated models on seen and unseen domains to shed light on the models’ generalization ability. Results are given in Table 2.

In Table 2, when comparing our baseline with pre-training strategies for each category (no curriculum), we note an average rise of 3.07, 19.6 and 10.97 respectively in BLEU when pre-training. Similar variations can be noted with METEOR, chrF++ and BLEURT metrics, albeit less striking. Based on BERT’s contextual embeddings, BLEURT gives a good estimate of the semantic correlation between prediction and references. Most of the time, n -gram-based and semantic metrics show perfect harmony. Top values of BLEURT are obtained for the same models than n -gram-based metrics. Therefore, all metrics tend to correlate proving a good agreement. The high improvements over unseen domains are easily explained due to the diversity of augmented data. New entities and domain-specific lexicon encountered better help to model out-of-distribution data relations. Thus, such results underline usefulness of external corpora and strengthen the need of pre-trained model

to lexicalize KBs.

On seen categories, our baseline give an acceptable BLEU score of 55.24. However, for out-of-domain generation, all models demonstrate severe shortcomings. Tested on unseen entities, our baseline shows a BLEU drop of 42.34 to reach 12.9 BLEU. We witness similar and even more important loss in unseen categories. In the case that predicates are unknown to the model, it is hard to generate consistent description of the input RDFs. From seen categories to unseen categories, our baseline is nearly penalized by a factor of 5. Such effect is tempered with our pre-training approaches. The average drop in BLEU of pre-trained models (without curriculum learning) is 25.6 from seen categories to unseen entities, and 35.5 from seen to unseen categories.

Unexpectedly, when a curriculum learning approach is used, we witness drops in performance. This is counter-intuitive and opposed to previous experiment on our validation set. CL seems to help lightly when the model was pre-trained with both external corpora WS1 and ST1. We propose to investigate the reason of such outcome in future work.

Interestingly, best results are revealed with a pre-training on ST1, exclusively. With 5 times less data, ST1 leads to better performance. The extracted triples surely include inaccurate triples. Notwithstanding the imperfect quality of the ST1 dataset, its use contributes to generalization ability. On the contrary, we report that denoising pre-training does not show satisfying results when combined with our pre-training on the ST1 dataset, leading eventually to negative effect. This may be caused by an input distribution discrepancy between WS1 and ST1 input. Denoising pre-training doesn’t require triples as input but a noisy sentence. Mismatch between this representation and the triples

Seen categories							
WebNLG	WS1	ST1	CL	BLEU	METEOR	chrf++	BLEURT
✓				55.24	0.401	0.680	0.56
✓	✓			57.3	0.417	0.701	0.58
✓		✓		59.32	0.428	<u>0.712</u>	<u>0.6</u>
✓	✓	✓		57.49	0.420	0.702	0.59
✓			✓	54.78	0.399	0.676	0.56
✓	✓		✓	56.94	0.417	0.701	0.58
✓		✓	✓	58.36	0.422	0.703	0.61
✓	✓	✓	✓	<u>58.81</u>	<u>0.427</u>	0.713	<u>0.6</u>

Unseen entities							
WebNLG	WS1	ST1	CL	BLEU	METEOR	chrf++	BLEURT
✓				12.9	0.167	0.319	-0.62
✓	✓			29.16	0.301	0.518	0.1
✓		✓		35.77	0.326	0.565	0.26
✓	✓	✓		32.33	0.310	0.535	0.18
✓			✓	11.94	0.156	0.295	-0.68
✓	✓		✓	28.83	0.295	0.509	0.05
✓		✓	✓	31.99	0.312	0.537	0.16
✓	✓	✓	✓	<u>32.69</u>	<u>0.315</u>	<u>0.542</u>	<u>0.19</u>

Unseen categories							
WebNLG	WS1	ST1	CL	BLEU	METEOR	chrf++	BLEURT
✓				11.17	0.162	0.310	-0.64
✓	✓			21.02	0.265	0.452	-0.06
✓		✓		<u>23.26</u>	0.288	0.485	0.03
✓	✓	✓		23.42	0.275	<u>0.469</u>	<u>-0.02</u>
✓			✓	12.45	0.156	0.298	-0.63
✓	✓		✓	21.84	0.263	0.451	-0.04
✓		✓	✓	22.84	0.273	0.465	-0.05
✓	✓	✓	✓	22.72	<u>0.276</u>	0.466	<u>-0.02</u>

Table 2: Ablation study: Automatic results on the official WebNLG test set for different learning strategy. For each experiment, we provide testing performance on seen categories (top), unseen entities (middle) and unseen categories (bottom). Bold and underlined values correspond to the best and second-best results respectively.

linearization may be the culprit of such side-effect.

As for the WebNLG 2020 challenge, participants were requested to submit a single model for evaluation. In our case, we decided to submit our Transformer pre-trained with both WS1 and ST1 (without curriculum). After an analysis of the generated text, the model trained with much more data has a tendency to be much more fluent and aggregate information better. [Castro-Ferreira et al. \(2020\)](#) report human evaluation based on different criteria: data coverage, relevance, correctness, text structure and fluency. For each criterion, a value between 0 (complete disagreement) and 100 (complete agreement) is given. The scores are normalized and

then clustered into groups such that models of a same cluster do not show any significant statistical differences in their scores⁸. When tested on seen categories, we note that our submitted model is competitive with other teams’ models. However, on unseen data, although a significant improvement over a simple training of a Transformer, our model shows limitations compared to other participants. A lower rank is systematically observed for both unseen entities and unseen categories. We assume that a delexicalization step and a much massive

⁸Details of the evaluation procedure is outlined in ([Castro-Ferreira et al., 2020](#)) and scores are publicly available at <https://gerbil-nlg.dice-research.org/gerbil/webnlg2020resultshumaneval>

pre-training as (Devlin et al., 2019; Radford et al., 2019) may help to improve our generalization ability on unseen domains. This direction is also to be explored in future work.

7 Conclusion

We proposed to use the Transformer model with data augmentation for the RDF-to-text task on the WebNLG dataset. We took advantage of Wikipedia to build new datasets for different pre-training strategies. We investigated the effect of denoising autoencoder objective and data augmentation as pre-training approaches. We studied the impact of curriculum learning while fine-tuning our model on WebNLG data. Further analysis demonstrated that pre-training is highly beneficial for knowledge verbalization. Although noisy, our massive data augmentation contributed to generate much accurate textual content. As future work, we will perform a deep analysis of the robustness of our models to unseen entities and categories.

8 Acknowledgment

This work was partially funded by the ANR Cifre conventions N°2020/0400 and N°2018/1559 and Orange-Labs Research. This work was granted access to the HPC resources of IDRIS under the allocation 2020-AD01101126 made by GENCI. We also thank Thibault Cordier and Imen Akermi for the fruitful discussions.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- 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.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. [Curriculum learning](#). In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, page 41–48, New York, NY, USA. Association for Computing Machinery.
- Thiago Castro-Ferreira, Claire Gardent, Nikolai Ilinykh, Chris van der Lee, Simon Mille, Diego Moussalem, and Anastasia Shimorina. 2020. The 2020 bilingual, bi-directional webnlg+ shared task: Overview and evaluation results (webnlg+ 2020). In *Proceedings of the 3rd WebNLG Workshop on Natural Language Generation from the Semantic Web (WebNLG+ 2020)*, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Philipp Cimiano, Janna Lüker, David Nagel, and Christina Unger. 2013. [Exploiting ontology lexica for generating natural language texts from RDF data](#). In *Proceedings of the 14th European Workshop on Natural Language Generation*, pages 10–19, Sofia, Bulgaria. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: pre-training text encoders as discriminators rather than generators](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- 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.
- Bhuvan Dhingra, Manaal Faruqui, Ankur P. Parikh, Ming-Wei Chang, Dipanjan Das, and William W. Cohen. 2019. [Handling divergent reference texts when evaluating table-to-text generation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4884–4895. Association for Computational Linguistics.
- Daniel Duma and Ewan Klein. 2013. [Generating natural language from linked data: Unsupervised template extraction](#). In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) – Long Papers*, pages 83–94, Potsdam, Germany. Association for Computational Linguistics.
- A. Gammerman, V. Vovk, and V. Vapnik. 1998. Learning by transduction. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, UAI'98*, page 148–155, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Hanning Gao, Lingfei Wu, Po Hu, and Fangli Xu. 2020. [Rdf-to-text generation with graph-augmented structural neural encoders](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3030–3036. International Joint Conferences on Artificial Intelligence Organization. Main track.

- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017a. [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.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017b. [The WebNLG challenge: Generating text from RDF data](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 124–133, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Qipeng Guo, Zhijing Jin, Xipeng Qiu, W. Zhang, D. Wipf, and Zheng Zhang. 2020. [Cyclegt: Unsupervised graph-to-text and text-to-graph generation via cycle training](#). *ArXiv*, abs/2006.04702.
- Hayate Iso, Chao Qiao, and Hang Li. 2020. [Fact-based Text Editing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 171–182, Online. Association for Computational Linguistics.
- Vladimir Iosifovich Levenshtein. 1965. Binary codes capable of correcting deletions, insertions, and reversals. In *Doklady Akademii Nauk*, volume 163, pages 845–848. Russian Academy of Sciences.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Diego Marcheggiani and Laura Perez-Beltrachini. 2018. [Deep graph convolutional encoders for structured data to text generation](#). In *Proceedings of the 11th International Conference on Natural Language Generation*, pages 1–9, Tilburg University, The Netherlands. Association for Computational Linguistics.
- David Mareček and Rudolf Rosa. 2019. [From balustrades to pierre vinken: Looking for syntax in transformer self-attentions](#). In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 263–275, Florence, Italy. Association for Computational Linguistics.
- Diego Moussallem, Dwaraknath Gnaneshwar, Thiago Castro Ferreira, and Axel-Cyrille Ngonga Ngomo. 2020. [Nabu - multilingual graph-based neural rdf verbalizer](#). In *The Semantic Web – ISWC 2020*. Springer International Publishing.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- 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.
- Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. 2019. [To tune or not to tune? adapting pre-trained representations to diverse tasks](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 7–14, Florence, Italy. Association for Computational Linguistics.
- Maja Popović. 2015. [chrF: character n-gram f-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Ehud Reiter and Robert Dale. 1997. [Building applied natural language generation systems](#). *Nat. Lang. Eng.*, 3(1):57–87.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. [Get to the point: Summarization with pointer-generator networks](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [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.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

- Xiantang Sun and Chris Mellish. 2007. [An experiment on free generation from single rdf triples](#). *Proceedings of the 11th European Workshop on Natural Language Generation, ENLG 07*, pages 105–108.
- Xiantang Sun and Christopher Stuart Mellish. 2006. Domain independent sentence generation from rdf representations for the semantic web. In *ECAI06 Combined Workshop on Language-Enabled Educational Technology and Development and Evaluation of Robust Spoken Dialogue Systems*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 3104–3112. Curran Associates, Inc.
- Bayu Distiawan Trisedya, Jianzhong Qi, Rui Zhang, and Wei Wang. 2018. [GTR-LSTM: A triple encoder for sentence generation from RDF data](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1627–1637, Melbourne, Australia. Association for Computational Linguistics.
- Bo-Hsiang Tseng, Jianpeng Cheng, Yimai Fang, and David Vandyke. 2020. [A generative model for joint natural language understanding and generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1795–1807, Online. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. 2019. [Xlnet: Generalized autoregressive pretraining for language understanding](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 5753–5763. Curran Associates, Inc.
- Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. 2019. [ERNIE: Enhanced language representation with informative entities](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451, Florence, Italy. Association for Computational Linguistics.
- Chao Zhao, Marilyn Walker, and Snigdha Chaturvedi. 2020. [Bridging the structural gap between encoding and decoding for data-to-text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2481–2491, Online. Association for Computational Linguistics.
- Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, and Maosong Sun. 2018. [Graph neural networks: A review of methods and applications](#). *CoRR*, abs/1812.08434.
- J. Zhu, T. Park, P. Isola, and A. A. Efros. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 2242–2251.
- Yaoming Zhu, Juncheng Wan, Zhiming Zhou, Liheng Chen, Lin Qiu, Weinan Zhang, Xin Jiang, and Yong Yu. 2019. [Triple-to-text: Converting rdf triples into high-quality natural languages via optimizing an inverse kl divergence](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19*, page 455–464, New York, NY, USA. Association for Computing Machinery.