

PARENTing via Model-Agnostic Reinforcement Learning to Correct Pathological Behaviors in Data-to-Text Generation

Clément Rebuffel^{1,2}, Laure Soulier¹, Geoffrey Scoutheeten², and Patrick Gallinari^{1,3}

¹LIP6, Sorbonne Université, France

²BNP Paribas, France

³Criteo AI Lab, Paris

firstname.lastname@{lip6.fr;bnpparibas.com}

Abstract

In language generation models conditioned by structured data, the classical training via maximum likelihood almost always leads models to pick up on dataset divergence (i.e., hallucinations or omissions), and to incorporate them erroneously in their own generations at inference. In this work, we build on top of previous Reinforcement Learning based approaches and show that a model-agnostic framework relying on the recently introduced PARENT metric is efficient at reducing both hallucinations and omissions. Evaluations on the widely used WikiBIO and WebNLG benchmarks demonstrate the effectiveness of this framework compared to state-of-the-art models.

1 Introduction

Data-to-Text aims at generating natural language descriptions from structured data (Reiter et al., 2005); fostered by recent advances on neural approaches and the emergence of large scale datasets made of (structured-data, reference text) pairs (Lebret et al., 2016; Gardent and Perez-Beltrachini, 2017; Wiseman et al., 2017). Figure 1 illustrates an example from the WikiBIO dataset (Lebret et al., 2016). These datasets are either hand-crafted via crowdworkers or automatically built by aligning sources found on the Internet. As such, reference texts might include divergences of two types, limiting the ability of generation models to produce realistic descriptions. First, reference texts might contain information not grounded in the source data; especially for automatically constructed datasets, where references were not written with the source-data description task in mind. For instance, the phrase “who served as lieutenant [...]” in Figure 1 has no basis in the associated infobox. Second, reference texts do not always cover the entirety of the table (items *Battles/wars* in Figure 1). In most settings, this second point is referred to as

content selection and is inherent of most data-to-text tasks. However, some hand-crafted datasets are designed where annotators are asked to transcribe every fields, with models also expected to do the same. In this case, incomplete references (i.e. where some part of the source data is missing from the realization) can lead to models failing to learn to transcribe all information, and only partially cover data-sources at inference.

Divergence in training examples leads to hallucinated/omitted content in model output; which is a well-known problem in neural approaches for text generation (Rohrbach et al., 2018). This problem arises both from the training procedure (training via maximum likelihood leads to language models strongly mimicking human behaviors), and from the testing protocols. Indeed, current standard metrics only measure similarity (such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Banerjee and Lavie, 2005)) to ground truth reference texts and do not fully capture relevance to the source data. Thus, there is no distinction between a mismatch caused by a paraphrase, poor lexicalization of content, or made-up/incorrect statement, leading to imperfect model selection. While a number of work argue for the need for novel automatic evaluation method (Reiter and Belz, 2009; Reiter, 2018; Novikova et al., 2017), to the best of our knowledge only Wiseman et al. (2017) and Dhingra et al. (2019) propose metrics based on both the reference and the source data.

Recently, different regularization methods have also been proposed to mitigate the negative influence of divergences in reference texts. These approaches can be either at the dataset level (Dušek et al., 2019), where authors propose techniques to clean/standardize instances; or at the training level (Tian et al., 2019), where authors propose novel neural modules designed to limit hallucinations/omissions. However, these

Sir James Love			
Born	1789	Battles/wars	Peninsular War Hundred Days
Died	13 January 1866	Awards	Knight Grand Cross of the Order of the Bath Knight of the Royal Guelphic Order
Allegiance	 United Kingdom		
Service/branch	 British Army		
Rank	General		

REFERENCE: general sir james **frederick** love gcb kh (1789 -- 13 january 1866) was a british army officer **who served as lieutenant governor of jersey** .

LSTM: general sir james love gcb (1789 -- 13 january 1866) was a british army officer **who went on to be commander-in-chief , the army** .

LSTM+RL: general sir james love gcb (1789 -- 13 january 1866) was a british army officer .

Figure 1: Example from WikiBIO where PARENTing (LSTML+RL) led to reduce hallucinated content. In yellow: divergent phrases in the reference; in red: hallucinated phrases from LSTM (without RL training) generation.

approaches are severely limited: e.g., they require significant annotation labor, model-specific tricks and/or manual tuning. Furthermore, virtually all proposed neural approaches still suffer from 1) *exposure bias* and 2) *inconsistency between train/test measurement*. Indeed, current neural models are trained via a mechanism called *teacher forcing* (Williams and Zipser, 1989), where the decoder is fed the previous correct token, no matter its actual prediction (1), in order to maximize the log-likelihood of the target sentence (including divergent phrases), but are evaluated through the previously discussed n-gram metrics (2). See Section 3.3 for a more detailed discussion about this subject.

To the best of our knowledge, there have been few approaches (Liu et al., 2019a,b) focused on the training procedure. Liu et al. (2019a) train a hierarchical encoder-decoder on three auxiliary tasks (namely *sequence labeling*, *text auto-encoder* and *multi-labeling classification*) which are meant to guide the decoding process. Closest to our work, Liu et al. (2019b) propose a novel neural module for constrained attention, along with a reinforcement learning (RL) training procedure based on BLEU and TFIDF. In our work, to remedy the above shortcomings and building upon the work of Liu et al. (2019b), we show that no novel neural module is necessary to handle hallucinations and omissions. We propose a model-agnostic RL framework, called PARENTing, where pretrained models are further trained with a self-critical policy gradient algorithm (Rennie et al., 2016) to limit the impact of divergences in training examples on text generation. Specifically, we use the PARENT metric (Dhingra et al., 2019) which exhibits a strong correlation with human evaluation, while being easier to use out of the box. We provide extensive automatic evaluations on two data-to-text model families (LSTMs and Transformers) on two widely

used benchmarks (WikiBIO and WebNLG), as well as a more focused human evaluation on WikiBIO. We report new state of the art PARENT scores on both datasets while BLEU scores are on par with previous SOTA approaches, which shows that our framework efficiently reduces pathological behaviors while keeping generation fluent.

2 Related Work

2.1 Text Generation from Structured Data

Data-to-text models can be classified in two broad categories: knowledge-based models and stochastic/data-driven approaches (Gatt and Krahrmer, 2018). The former approaches (Reiter and Dale, 2000) are driven by experts’ knowledge, leading to a pipeline architecture split into subtasks: content selection and text structuring (*macroplanning*), sentence planning (*microplanning*) and generating actual sentences (*surface realisation*). While accurate and efficient at inference time, these methods require significant manual efforts for new use-cases. In contrast, data-driven approaches tend to blur the distinction between these subtasks with end-to-end training on large corpora of aligned input data and output text (Gatt and Krahrmer, 2018). End-to-end methods have been proposed early, such as (Chen and Mooney, 2008) who apply statistical machine translation techniques to the sport-casting domain. Recent neural approaches now propose to leverage progress in deep learning to represent these data into a semantic vector space (also called embedding space) and stem from the neural machine translation domain (Lebret et al., 2016; Puduppully et al., 2019; Wiseman et al., 2017). Particularly, Wiseman et al. (2017) propose the now by default back-bone data-to-text architecture, with an attention mechanism (Bahdanau et al., 2014), which computes a context focused on important elements from the input, and a copy mechanism (Gulcehre et al., 2016; See et al., 2017) to deal with unknown or rare words.

To address domain-specific constraints, a common approach is to build architectures that explicitly model the key-value structure of the input table (Nie et al., 2018; Liu et al., 2018, 2019a,b; Rebuffel et al., 2020). Additional work (Puduppully et al., 2019) introduces dynamic encoding updating, where the model updates part of the source data encoding at each decoding step in order to accurately guide the decoder throughout generation.

While these models produce fluent and domain-comprehensive outputs, several pathological behaviors have been identified, echoing similar issues in other text generation tasks (e.g. in image captioning (Rohrbach et al., 2018) or in summarization (Kryściński et al., 2019)).

2.2 Pathological Hallucinations in Data-to-Text

Training neural model on data-to-text tasks requires large corpora (Lebret et al., 2016; Novikova et al., 2017; Gardent and Perez-Beltrachini, 2017; Wiseman et al., 2017). Different pathological behaviors arise from the datasets, depending on the methodology underlying their construction. First, for hand-crafted datasets (Novikova et al., 2017; Gardent and Perez-Beltrachini, 2017), crowdworkers sometimes fail to cover all information from the data source in reference text. Second, automatically constructed datasets from possibly different internet sources do not guarantee data sources and texts to be aligned completely. Both these limitations induce neural generation model to omit information in the first case or suffer from hallucinations (i.e., they mistakenly learn to generate ungrounded/false statements) in the other.

To deal with these pathologies, previous work operate either at the dataset level, or at the training level. At the dataset level, Dušek et al. (2019) show that cleaned data can significantly improve system ability to produce fact-accurate text. In a different direction, Nie et al. (2019) apply a method similar to knowledge distillation (Hinton et al., 2015): they train a Natural Language Understanding module to reconstruct tables from text references and show that a vanilla sequence-to-sequence model trained on the refined data has improved content correctness in both human and automatic evaluations. At the training level, Wiseman et al. (2017), for instance, propose to include a reconstruction loss aiming at reconstructing the source table from the hidden states of the decoder. In an other direction,

Perez-Beltrachini and Lapata (2018) propose a classifying neural network, trained (using a manually annotated dataset) to label text tokens depending on their alignment with the associated table. They use these labels in an RL framework to generate sentences with a maximum of aligned tokens. However these approaches are either costly in human labor or specific to hand-crafted datasets where the input data matches exactly the reference texts (thus deal with omissions but not hallucinations). Indeed reconstruction tasks are not compatible with the content selection subtask of Data-to-Text.

Proposing both a novel coverage-constrained attention and a BLEU/TFIDF-based reward, (Liu et al., 2019b) constitutes a first approach to a model-agnostic framework. However their proposed coverage is still task specific (and goes against content-selection): while they increase the state-of-the-art BLEU on WikiBIO, they underperform encoder-decoder models on the PARENT benchmark.

Until recently, the NLG research community cruelly lacked ways to automatically evaluate model outputs. Despite work on effective human evaluation (Amidei et al., 2019), and on the need for better automated metrics (Reiter and Belz, 2009; Novikova et al., 2017), to the best of our knowledge, only Wiseman et al. (2017) and Dhingra et al. (2019) recently proposed improvement over the widely used BLEU. Wiseman et al. (2017) propose to use an auxiliary neural model, trained to extract structured records from the generated text for evaluation. Two texts can then be compared through their sequences of extracted records. This information retrieval-based approach suffers from domain specificity, as the released model only works in the closed-domain of basketball journalistic summaries, and requires precise tagging of gold references which can be impossible to provide in most settings. Furthermore, Dhingra et al. (2019) propose a new metric PARENT, and show that this metric strongly correlates with human annotators and can replace previous n-gram- and information retrieval-based metrics.

Our contribution differs from previous work in several aspects. First, our proposed framework is model-agnostic and can be used with any neural model. Second, instead of focusing on only one domain and/or one issue (e.g., omissions in hand-crafted datasets or hallucinations in automatically constructed datasets), it is setting agnostic and tackles both hallucinations and omissions at once by

leveraging the PARENT F-score (which combines precision and coverage against the source data). Finally, no manual preprocessing or pre-tagging is required: models are trained via a flexible training protocol and distantiate themselves from faulty training examples.

3 Model-Agnostic Reinforcement Learning for Reducing Divergences

We propose PARENTing, a model-agnostic RL framework for data-to-text aiming at reducing divergences. It is based on the self-critical policy gradient algorithm (Paulus et al., 2018) and leverages the PARENT metric (Dhingra et al., 2019).

3.1 Background and Notations

Notations. We consider the general setting of data-to-text and the notations introduced by Dhingra et al. (2019). Let us consider a dataset of J pairs (structured-data, reference), denoted $\mathcal{D} = \{(T_j, R_j)\}_{j=1}^J$ where:

- $T := \{r_k\}_{k=1}^K$ is a collection of K records (*entity, attribute, value*), where K is variable among instances;

- $R := [y_1^*, \dots, y_L^*]$ is the reference text associated to T , composed of L tokens y^* , where L is variable among instances;

We also consider a data-to-text neural model, denoted f_θ (where θ are the model parameters), pre-trained to maximize the likelihood of the reference, via teacher forcing (Williams and Zipser, 1989):

$$\mathcal{L}_{ml} = - \sum_{t=1}^L \log f(y_t^* | y_{t-1}^*, \dots, y_1^*, T, \theta) \quad (1)$$

PARENT metric. PARENT (Precision And Recall of Entailed N-grams from the Table) (Dhingra et al., 2019) aims at evaluating the precision and recall/coverage of a candidate generation G given the (source table, reference) pair (T, R) via n-grams ($n = 1, \dots, 4$) comparison. This metric is divided in three scores:

- Entailed precision \mathbf{E}_p is the fraction of n-grams from G which are either found in R or T ;

- Entailed recall/coverage \mathbf{E}_r . Recall $E_r(R)$ is the fraction of n-grams from $R \cap T$ which are found in G ; Coverage $E_r(T)$ is the fraction of n-grams from T which are found in G . Recall and coverage are combined using a geometric average:

$$\mathbf{E}_r := E_r(R)^\lambda E_r(T)^{1-\lambda} \quad (2)$$

- F-score: combination of precision and recall.

3.2 Overview and Research Objectives

Our framework for reducing pathological behaviors is based on the following research objectives:

- O1: the framework should be generic and should work with any neural model;

- O2: the model should try and distantiate itself from the reference enough to stop mimicking problematic behaviors;

- O3: by combining precision, recall and coverage, PARENT is a good proxy for human assessment of a candidate text against its source data and reference (Dhingra et al., 2019; Tian et al., 2019);

- O4: discrete metrics can be gamed to artificially increase the score while not gaining in readability or relevance (Liu et al., 2016).

Therefore, we propose a training protocol similar to (Paulus et al., 2018), with a mixed objective function combining the standard maximum-likelihood loss \mathcal{L}_{ml} with a custom reinforcement loss \mathcal{L}_{rl} . This ensures that models do not lose fluency by gaming the discrete metric (objective O4):

$$\mathcal{L} := \gamma \mathcal{L}_{rl} + (1 - \gamma) \mathcal{L}_{ml} \quad (3)$$

where γ is a weight factor. We note that O1 is satisfied, as this loss function \mathcal{L}_{rl} can be applied to train any neural model f_θ . In what follows, we give a description of the proposed reinforcement learning framework.

3.3 PARENTing: Self-critical Gradient Policy Learning

Numerous work (Wu et al., 2016; Rennie et al., 2016) have outlined that training via teacher-forcing (maximizing the log-likelihood of reference texts) does not always produce the best results on evaluation metrics. This is in part due to *exposure bias* (Ranzato et al., 2016), where models are trained using the true gold sequence during training and are never exposed to their possible mistakes.

We therefore propose to alter the standard rigid training protocol and further train models via reinforcement learning as a counter-measure to these issues, where models can now learn a more flexible policy based on a metric more representative of human judgment, satisfying objective O2. Following objective O3, we shape our reward around the PARENT metric which has been shown to strongly correlate with human judgement in term of precision and recall of a generated text against a source table and a reference. Models are somewhat over-fitted to our training set due to pretraining, and are

hence at risk of earning high rewards on easy examples (i.e. with faithful reference targets) and poor rewards on hard examples (i.e. with divergent reference targets). To deal with this issue and ensure that the reward reflects the actual improvement made over the pretraining, we propose to follow a growing body of work in text summarization (Paulus et al., 2018; Scialom et al., 2019) and apply the self-critical policy gradient training protocol (Renzie et al., 2016), using the REINFORCE (Williams and Peng, 1991) algorithm.

More particularly, models are now sampled using their Markov property (that is one token at a time, and computing the next distribution given the previous chosen token). A first candidate sequence Y^c is randomly sampled following the outputted distribution. A second baseline sequence Y^b is generated, this time via greedy decoding (mimicking beam search generation during inference, with a beam of size 1). This baseline sequence acts as a difficulty proxy of the current training instance. The reward given to the candidate sequence is the improvement in PARENT score it brings over the baseline sequence:

$$r(Y^c) = \text{PARENT}(Y^c) - \text{PARENT}(Y^b) \quad (4)$$

Finally, the loss to be minimized during this part of training is:

$$\mathcal{L}_{rl} = -r(Y^c) \sum_{t=1}^L \log f(y_t^c | y_{t-1}^c, \dots, y_1^c, T, \theta) \quad (5)$$

Minimizing Equation 5 leads to increase reward expectation. Indeed, we maximize the conditional likelihood of the candidate sequence Y^c when it obtains a higher reward than the baseline sequence Y^b , or on the contrary we decrease its likelihood in case of a lower reward.

4 Experimental setup

4.1 Data-to-text benchmarks

WikiBIO (Lebret et al., 2016) This dataset contains 728, 321 infoboxes, automatically paired with the first sentence of the corresponding article of the English Wikipedia. We follow the data partition introduced with the dataset which yields 80% of all instances for the training set, 10% for the development set and 10% for the evaluation set. Reference texts are of average length 26 words while infoboxes have on average 12 non-empty fields. This

dataset has been built automatically from sources that were not meant for a text-generation task and contains a significant amount of divergence between the source data and the target descriptions (62% of the references mention extra information not grounded in the infobox (Dhingra et al., 2019)).

WebNLG (Gardent and Perez-Beltrachini, 2017)

This dataset contains 35, 970 sets of RDF records mapped to natural language descriptions. Each set has up to 7 records, and one or more gold references of average size 22 words. We follow the partition introduced with the dataset, which yields 1612/1619 instances as a development/evaluation set. This dataset has been hand-crafted specifically for the task of surface realization and systems are expected to summarize all records. Note that here we compare ourselves on the *seen* partition, where every attribute is seen during training (however, entities and values can be new).

4.2 Evaluation metrics

We evaluate our approach using both automated metrics and human judgment. We report BLEU scores (Papineni et al., 2002) as well as PARENT (precision, recall and F1) scores (Dhingra et al., 2019). For all scores higher is better. While BLEU is the historical metric in all text generation tasks, PARENT scores have a significantly stronger correlation with human evaluators (Dhingra et al., 2019) (0.478 vs. 0.913 for BLEU and PARENT resp.).

We perform qualitative evaluation following the best practices outlined by (van der Lee et al., 2019). Our human annotators are males and females from several countries across Europe, between 20 and 55 years old and proficient in English. Annotators are shown a randomly selected table, together with the corresponding descriptions, both from the dataset and the models that are being evaluated. Annotators are asked, for each sentence, to score its fluency (as *Fluent*, *Mostly fluent*, or *Not fluent*) factualness (likewise), and coverage (in terms of the number of realized rows). Sentences are shuffled to avoid any bias. Following Tian et al. (2019), we first tasked three expert annotators to annotate a pilot batch of 50 sentences. Once assured all Inter-Annotator Agreements were approx. 78%, we asked several annotators to annotate an additional sentence sample to reach 100 instances (where each instance consists of one table and three associated outputs).

Model		WikiBIO				WebNLG			
		BLEU	PARENT (Prec. / Rec. / F1)			BLEU	PARENT (Prec. / Rec. / F1)		
Baselines	S2S+FA+RL	45.49	76.1	45.9	54.8	-	-	-	-
	Confident PG	38.10	79.52	40.60	51.38	-	-	-	-
	GCN	-	-	-	-	55.9	-	-	-
	Gardent-LSTM	-	-	-	-	54.03	-	-	-
Our models	LSTM	42.80	78.70	45.16	55.10	54.9	73.15	69.91	69.67
	LSTM+RL	44.17*	80.01*	46.60*	56.72*	63.20*	74.51	70.68*	71.27*
	Transformer	41.09	80.02	44.31	54.74	53.45	75.17	64.15	67.38
	Transformer+RL	42.40	80.37	45.83*	56.15*	57.07*	75.06	66.70*	69.01*

Table 1: Evaluation on WikiBIO and WebNLG. *: p-value < 0.001 of Student T-test comparing with/without RL.

4.3 Scenarios and Baselines

We measure the impact of our framework on two families of models:

- *LSTMs*. Our implementation of (See et al., 2017). It is the back-bone data-to-text model based on a bi-LSTM with attention mechanism and augmented with a conditional copy mechanism to deal with rare or unseen words.
- *Transformers*. Our implementation of (Vaswani et al., 2017), the transformer encoder-decoder, augmented with a conditional copy mechanism.

These models are denoted **LSTM** or **Transformer** when trained via maximum likelihood and **LSTM+RL** or **Transformer+RL** when further trained using the PARENTing framework.

We also report SOTA models for each dataset respectively (i.e., achieving the strongest score either BLEU or PARENT) :

- For WikiBIO, we report the BLEU and PARENT scores of two baselines: 1) **S2S+FA+RL** (Liu et al., 2019b) which uses a standard encoder-decoder structure, with an attention mechanism constrained to cover all table attributes and an RL training procedure with a reward shaped by BLEU and TFIDF; 2) **Confident PG** (Tian et al., 2019): a neural module which assigns a confidence score to each output words, and trims the generated sequence from any word below a specified threshold. They report higher precision but lower fluency.
- For WebNLG, we report the BLEU score of **GCN Marcheggiani and Perez-Beltrachini (2018)**. They propose a graph convolutional network which explicitly models the structure of graph-like data. For additional context, we also report a baseline score introduced by the original paper **Gardent and Perez-Beltrachini (2017)**. The used model **Gardent-LSTM** is the same as our scenario LSTM.

4.4 Implementations details

We describe here key implementation details (other details needed for reproducibility will be given alongside the code if accepted). We set the λ of Equation (2) to 1 during training. This was done 1) because coverage is against the content selection task on most data-to-text tasks 2) to reduce the computing cost, as coverage is obtained by computing Longest Common Subsequence for all n-grams contained in the table. We note however that we kept $\lambda = 0.5$ for evaluation following Dhingra et al. (2019); Tian et al. (2019). Preliminary experiment on γ from Equation 3 showed that the initial value of 0.9987 proposed by Paulus et al. (2018) was not satisfying: fluency dropped drastically and while models obtained significantly higher PARENT scores, BLEU score was at less than half what previous models were able to achieve. We therefore used a more conservative value of $\gamma = 0.9$. Inputs were fed to the neural networks following Lebret et al. (2016): each word is represented as a 4-tuple (value, field, p+, p-) where p+ (resp. p-) is the position (resp. reverse position) of value in field. For example, the line (Name, Barack Obama) is presented as [(Name, Barack 1, 2), (Name, Obama, 2, 1)]. In WebNLG, where tables include several entities, a 5th element was introduced for entity index, as well as tokens for the entities’ names. Models are first trained via maximum likelihood training. We select the best performing checkpoint given a development set and start the mixed-objective training from there. We implemented our framework using OpenNMT (Klein et al., 2017). Data and code are available online: <https://github.com/KaijuML/PARENTing-rl>

5 Results

Table 1 summarizes the BLEU and PARENT scores obtained by the baselines and our scenar-

ios on WikiBIO and WebNLG benchmarks. Please note that while no previous work report PARENT scores on WebNLG, our scenario *LSTM* is a re-implementation of the baseline *Gardent-LSTM*. The obtained and reported BLEU scores being very close, we can consider that our PARENT scores are also the ones of *Gardent-LSTM*.

From a general point of view, We can see that our PARENTed models *LSTM+RL* and *Transformer+RL* obtain generally higher BLEU and PARENT metrics over all scenarios and baselines — except the BLEU score for WikiBIO. More particularly, we can outline the following statements.

- The comparison of our scenarios (without/with our PARENTing framework) outlines increases in score ranging from +1.6% to +3.4% on WikiBIO, and from +1.1% to +15% on WebNLG; with significant improvements in 12/16 comparison settings. This suggests that PARENTed models learn to describe source data with more precision (reduced hallucinations) and with greater details (increased recall/coverage), as can be seen in Figure 1.
- BLEU scores are on par with baselines on WikiBIO, and significantly better than the strongest model on WebNLG. Despite starting close to the baseline in terms of BLEU, PARENTing our models leads to new state of the art BLEU of 63.20, compared to a previous 55.9 for GCN, representing a 13% relative increase. This shows that models learn to lexicalize content more adequately than through maximum-likelihood training.
- More importantly, PARENTed models overperform all baselines on PARENT scores: on WikiBIO, our reinforced *LSTM* model achieves state of the art performance on F-score, increasing the previous score by 7% (from 52.81 to 56.72). PARENTed models also achieve better precision than both *S2S+FA+RL* and *Confident-PG*: 80.01/80.37 (for *LSTM+RL* and *Transformer+RL* resp.) against respectively 76.1 and 79.52. Contrasting with *Confident-PG* which sacrifices fluency for faithfulness, our models referring more precisely to information from the table (*i.e.* reducing hallucinations) does not come at the cost of less fluent output, as reflected by the BLEU scores of 44.17/42.40 against 38.1 for *Confident-PG*. We note that beyond precision, PARENTed models are trained simultaneously to retrieve more relevant information (*i.e.* reducing omissions), as they achieve SotA recall performance (46.60/45.83), again contrasting with *Confident-PG* which loses more than 11% to

	LSTM scenario
Avg length without RL	19.17
Avg length with RL	19.78
Avg length Δ	0.61
Δ of F-measure scores	1.62
Correlation	0.21* ($p < 0.001$)

Table 2: Comparative statistics between *LSTM* and *LSTM+RL* on WikiBIO. Δ for variation scores.

S2S+FA+RL, with 40.6 against 44.02.

- Altogether, the previous statements assess the model-agnosticity of our framework, as both model family (*LSTM* and *Transformer*-based scenarios) showed improvements on both datasets when fine-tuned with our PARENTing framework.
- We observe that our pre-trained scenarios (*LSTM* and *Transformer*) generally obtain higher results on WebNLG than on WikiBIO. This is due to the nature of datasets: WebNLG is hand-crafted with the explicit goal of full transcription of tables while WikiBIO is build automatically without rigorous alignment of data sources and reference texts. Despite some inevitable divergences, WebNLG is thus less noisy than WikiBIO.

Qualitative Evaluations. We aim to provide insight into what our PARENTing framework brings to models. Specifically, one might assume that a model could trivially learn to shorten output in order to increase precision, or on the contrary, to increase generation length to easily increase coverage by mechanically quoting tables more. We therefore 1) check for the framework impact on global generation length; 2) provide a more detail analysis on length distribution vs. effectiveness.

In this section, we focus exclusively on WikiBIO as it is the most challenging setting (larger vocabulary, more noisy, and content selection needed to generate biographies). To make results more readable, we focus on (*LSTM*, *LSTM+RL*) models.

We first report in Table 2 comparative statistics of generation length and score variations. We first note that on WikiBIO, there is no significant changes in sentence lengths after RL training (19.17 vs 19.78). We further find a correlation of 0.2 between the variation in length and the one in PARENT F-score which does not allow to conclude that longer texts lead to better scores. This suggests that the RL-trained model seems to have better performance even when generating shorter texts.

We investigate further the impact of PARENTing

	LSTM		LSTM+RL	
	Short	Long	Short	Long
Precision	78.75	77.76*	80.04	80.03
Recall	44.82	52.01*	46.27	52.41*
F-score	54.82	60.82*	56.44	61.57*
Nb-copy	9.25	11.51*	9.94	12.05*

Table 3: Effectiveness analysis depending on generation size. Nb-copy: average number of words copied from the source table. *: p-value < 0.001 for Student T-test comparing short/long generations.

on length distribution, and its influence on hallucinations/omissions (respectively measured by precision/recall). To do so, generated texts are splitted in two broad categories, short and long, using a KMeans algorithm calibrated on the length of human references. We exhibit two clusters, texts below/above 30 words, and compute PARENT scores conditioned on these clusters (see Table 3).

Considering hallucinations, where being precise should naturally be increasingly harder with sentence length, we first observe that for the pretrained model, precision tends to decrease with longer generation (78.75 vs. 77.76), while the RL-trained model is more robust and has constant precision independently of generation length¹.

Regarding omissions, Table 3 shows that while both scenarios have similar recall/coverage scores in long generations, the RL-trained model is more relevant; achieving a recall of 46.27 on short generations against 44.82 for the pretrained model. Finally, this leads to overall higher PARENT score for the PARENTed model than the pretrained model, independently of text lengths, showing that the RL-trained model choses more accurately when to stop generating early and when to pursue longer generation, adding additional information from the table to improve coverage.

Interestingly, we find that no matter the generation size, the PARENTed model shows more reliance on the source table, as it tends to directly copy words more often than its pretrained version.

Human Evaluations. To better measure subtleties which would not be captured by automatic metric in model outputs, we report human ratings in Table 4. Due to cost of human evaluation, we focus on the WikiBIO dataset, and three settings: the best model from the literature S2S+FA+RL, our model

¹This behavior, stopping generation early to avoid hallucinated statements, is illustrated in Figure 1.

Model	Fluency	Factualness	Coverage
Gold	91.7%	32.1%	4.25
S2S+FA+RL	84.5%	58.3%	4.45
LSTM+RL	94%	72.6%	4.22

Table 4: Results of the human evaluation on WikiBIO. The Fluency column reports the count of sentences labeled as “fluent” or “mostly fluent”.

LSTM+RL and gold sentences. It is worth noting that our results align with [Dhingra et al. \(2019\)](#), as we found that around two thirds of gold references contain divergences from their associated tables.

- The fluency scores highlight the need for a mixed objective loss, which leverages the MLE objective ability to produce fluent output, whereas RL alone (S2S+FA+RL) leads to less fluent output due to the discrete metric being used as a reward. Indeed, S2S+FA+RL obtains only a score 84%, compared to gold standard of nearly 92%, or our model’s score of 94%².

- Factualness scores show that both approaches greatly improve factualness over gold standard. However, S2S+FA+RL still lags behind our proposed approach, which is able to leverage the PARENT-based reward to constrain the system better than would the FA module.

- In contrast to factualness, S2S+FA+RL obtains better coverage performance than our approach, with 4.45 vs 4.25, showing that a component dedicated to coverage (either the FA module or the TFIDF part of the reward) leads to global outputs. Despite this, coverage is on par with gold standard.

6 Discussion and Conclusion

In this work, we have proposed a model-agnostic reinforcement learning framework for data-to-text aimed at reducing hallucinations and improving recall/coverage of relevant information. We shaped the reward based on PARENT ([Dhingra et al., 2019](#)), which is a recently proposed metric with a high correlation with human judgement.

This allows for a more flexible training, where the model learns to depend less on the reference and more on the source data. Framework effectiveness is assessed via thorough experiments on two model family (RNNs and Transformers) and two benchmarks (WikiBIO and WebNLG). Furthermore, quantitative and qualitative evaluations show

²Gold standard not being 100% is explained by preprocessing choices made at dataset creation (e.g. non-english languages are not always correctly transcribed).

that our PARENTing framework obtains better results than a dedicated attention module or a less source-relying reward.

However, this approach relies on the metric employed and crafting an effective metric is still an open problem. In particular, PARENT is designed for single-entity datasets, like WikiBIO and WebNLG, which is not reliable for more complex datasets containing multiple entities (i.g., the Ro-toWire dataset (Wiseman et al., 2017)). In this setting, the sentence “James Harden scored 20 points.” could achieve a high PARENT score if any player had scored 20 points in the game. An interesting future work would be the design of an evaluation metric more robust to dataset peculiarities.

Acknowledgements

We would like to thank the H2020 project AI4EU (825619) and the ANR JCJC SESAMS (ANR-18-CE23-0001) for supporting this work.

References

Jacopo Amidei, Paul Piwek, and Alistair Willis. 2019. The use of rating and likert scales in natural language generation human evaluation tasks: A review and some recommendations. In *Proceedings of the 12th International Conference on Natural Language Generation*.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. Accepted at ICLR 2015 as oral presentation.

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*. Association for Computational Linguistics.

David L. Chen and Raymond J. Mooney. 2008. Learning to sportscast: A test of grounded language acquisition. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*. ACM.

Bhuwan Dhingra, Manaal Faruqui, Ankur Parikh, Ming-Wei Chang, Dipanjan Das, and William Cohen. 2019. Handling divergent reference texts when evaluating table-to-text generation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics.

Ondřej Dušek, David M Howcroft, and Verena Rieser. 2019. Semantic Noise Matters for Neural Natural Language Generation. page 6.

Claire Gardent and Laura Perez-Beltrachini. 2017. A statistical, grammar-based approach to microplanning. *Computational Linguistics*, 43(1):1–30.

Albert Gatt and Emiel Krahmer. 2018. Survey of the state of the art in natural language generation: Core tasks, applications and evaluation. *J. Artif. Int. Res.*, 61(1):65–170.

Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. Pointing the unknown words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.

Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531.

Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. OpenNMT: Open-source toolkit for neural machine translation. In *Proc. ACL*.

Wojciech Kryściński, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Evaluating the factual consistency of abstractive text summarization.

Rémi Lebret, David Grangier, and Michael Auli. 2016. Neural text generation from structured data with application to the biography domain. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1203–1213, Austin, Texas. Association for Computational Linguistics.

Chris van der Lee, Albert Gatt, Emiel van Miltenburg, Sander Wubben, and Emiel Krahmer. 2019. Best practices for the human evaluation of automatically generated text. In *Proceedings of the 12th International Conference on Natural Language Generation, INLG 2019, Tokyo, Japan, October 29 - November 1, 2019*, pages 355–368.

Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Chia-Wei Liu, Ryan Lowe, Iulian Serban, Mike Noseworthy, Laurent Charlin, and Joelle Pineau. 2016. How NOT to evaluate your dialogue system: An empirical study of unsupervised evaluation metrics for dialogue response generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2122–2132, Austin, Texas. Association for Computational Linguistics.

Tianyu Liu, Fuli Luo, Qiaolin Xia, Shuming Ma, Baobao Chang, and Zhifang Sui. 2019a. Hierarchical encoder with auxiliary supervision for neural table-to-text generation: Learning better representation for tables. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33:6786–6793.

- Tianyu Liu, Fuli Luo, Pengcheng Yang, Wei Wu, Baobao Chang, and Zhifang Sui. 2019b. [Towards comprehensive description generation from factual attribute-value tables](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019*, pages 5985–5996.
- Tianyu Liu, Kexiang Wang, Lei Sha, Baobao Chang, and Zhifang Sui. 2018. [Table-to-text Generation by Structure-aware Seq2seq Learning](#). In *AAAI*.
- 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. Association for Computational Linguistics.
- Feng Nie, Jinpeng Wang, Jin-Ge Yao, Rong Pan, and Chin-Yew Lin. 2018. [Operation-guided neural networks for high fidelity data-to-text generation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3879–3889.
- Feng Nie, Jin-Ge Yao, Jinpeng Wang, Rong Pan, and Chin-Yew Lin. 2019. [A simple recipe towards reducing hallucination in neural surface realisation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2673–2679. Association for Computational Linguistics.
- Jekaterina Novikova, Ondřej Dušek, Amanda Cercas Curry, and Verena Rieser. 2017. [Why we need new evaluation metrics for NLG](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2241–2252. 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 on Association for Computational Linguistics, ACL '02*, pages 311–318, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2018. [A deep reinforced model for abstractive summarization](#). In *6th International Conference on Learning Representations, ICLR 2018*.
- Laura Perez-Beltrachini and Mirella Lapata. 2018. [Bootstrapping generators from noisy data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018*, pages 1516–1527. Association for Computational Linguistics.
- Ratish Puduppully, Li Dong, and Mirella Lapata. 2019. [Data-to-text generation with entity modeling](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019*, pages 2023–2035.
- Marc’Aurelio Ranzato, Sumit Chopra, Michael Auli, and Wojciech Zaremba. 2016. [Sequence level training with recurrent neural networks](#). In *4th International Conference on Learning Representations, ICLR 2016, San Juan, Puerto Rico, May 2-4, 2016, Conference Track Proceedings*.
- Clément Rebuffel, Laure Soulier, Geoffrey Scoutheeten, and Patrick Gallinari. 2020. [A hierarchical model for data-to-text generation](#). In *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part I*, pages 65–80.
- Ehud Reiter. 2018. [A structured review of the validity of bleu](#). *Computational Linguistics*, 44(3):393–401.
- Ehud Reiter and Anja Belz. 2009. [An investigation into the validity of some metrics for automatically evaluating natural language generation systems](#). *Computational Linguistics*, 35(4):529–558.
- Ehud Reiter and Robert Dale. 2000. *Building Natural Language Generation Systems*.
- Ehud Reiter, Somayajulu Sripada, Jim Hunter, Jin Yu, and Ian Davy. 2005. [Choosing words in computer-generated weather forecasts](#). *Artif. Intell.*, 167(1-2):137–169.
- Steven J. Rennie, Etienne Marcheret, Youssef Mroueh, Jerret Ross, and Vaibhava Goel. 2016. [Self-critical sequence training for image captioning](#). *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1179–1195.
- Anna Rohrbach, Lisa Anne Hendricks, Kaylee Burns, Trevor Darrell, and Kate Saenko. 2018. [Object hallucination in image captioning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4035–4045. Association for Computational Linguistics.
- Thomas Scialom, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2019. [Answers unite! unsupervised metrics for reinforced summarization models](#). 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 3244–3254. Association for Computational Linguistics.
- 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*, pages 1073–1083. Association for Computational Linguistics.
- Ran Tian, Shashi Narayan, Thibault Sellam, and Ankur P. Parikh. 2019. [Sticking to the facts: Confident decoding for faithful data-to-text generation](#). *CoRR*, abs/1910.08684.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17.

R. J. Williams and D. Zipser. 1989. [A learning algorithm for continually running fully recurrent neural networks](#). *Neural Computation*, 1(2):270–280.

Ronald J. Williams and Jing Peng. 1991. Function optimization using connectionist reinforcement learning algorithms.

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. Association for Computational Linguistics.

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google's neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.