

DATScore: Evaluating Translation with Data Augmented Translations

Abstract

The rapid development of large pretrained language models not only has revolutionized the field of Natural Language Generation (NLG), but also its evaluation. Inspired by the recent work of BARTScore: a metric leveraging the BART language model to evaluate the quality of generated text from various aspects, we introduce DATScore. DATScore uses data augmentation techniques to improve the evaluation of machine translation. Our main finding is that introducing data augmented translations of the source and reference texts is greatly helpful in evaluating the quality of the generated translation. We also propose two novel score averaging and term weighting strategies to improve the original score computing process of BARTScore. Experimental results on WMT show that DATScore correlates better with human meta-evaluations than the other recent state-of-the-art metrics, especially for low resource languages. Ablation studies demonstrate the value added by our new scoring strategies. Moreover, we report the performance of DATScore on 3 other NLG tasks than translation in our extended experiments. Code is publicly available¹.

1 Introduction

Massive pretrained language models have brought significant improvement to NLG tasks (Lewis et al., 2020). Recent systems can even generate texts of higher quality than human annotated ones (Peyrard, 2019). At the same time, standard metrics, such as BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004), for translation and summarization respectively, have not evolved for the past two decades (Bhandari et al., 2020). These metrics rely on surface lexicographic matches, making them particularly unsuitable for evaluating modern systems operating with embeddings at the semantic level, that often generate paraphrases (Ng and Abrecht,

2015). To address this issue, many metrics have been proposed (Sai et al., 2022), but none of them were widely adopted until the release of BERTScore (Zhang et al., 2019) and MoverScore (Zhao et al., 2019). These metrics take advantage of large pretrained language models like BERT (Devlin et al., 2019), which are now being used in nearly all NLP tasks (Qiu et al., 2020; Min et al., 2021).

In this work, we focus on the task of evaluating machine translation. We propose an extension of BARTScore (Yuan et al., 2021), a recent metric exploiting the BART seq2seq language model (Lewis et al., 2020) to evaluate the quality of generated text from various aspects. BARTScore covers four evaluation facets: Faithfulness, Precision, Recall, and F-score, derived from different generation directions between the *source* text, the *hypothesis* (the text generated by a system given the source), and the *reference* (the reference text for the generation, often provided by human annotators). The scores are obtained by pairing the 3 entities differently at the input or the output side of a trained seq2seq model for fetching conditional generation probabilities.

On the basis of BARTScore, and motivated by the general idea and positive effect of data augmentation techniques, we found that adding augmented, translated copies of the source and reference texts in BARTScore, can greatly help in evaluating the quality of the hypothesis translation. We also propose two novel score averaging and term weighting strategies to improve the original score computing process of BARTScore. Results and ablation studies show that our metric DATScore (Data Augmented Translation Score) outperforms the other recent state-of-the-art metrics, and our new scoring strategies are effective. Moreover, the performance of DATScore is also reported on 3 other NLG tasks than translation: data-to-text, summarization, and image captioning.

To the best of our knowledge, no prior work has

¹link will be provided upon acceptance.

been done on leveraging data augmentation techniques for untrained NLG evaluation metrics. Our work will help filling this gap. Our contributions include:

- 1) Inspired by BARTScore, we developed DATScore that incorporates augmented data translated from the source and reference texts. DATScore is an untrained and unsupervised translation evaluation metric, that offers a larger performance boost in evaluating low resource language generation. In contrast to other widely adopted metrics, DATScore can efficiently incorporate both the source and reference texts in the evaluation.
- 2) We introduced a novel one-vs-rest method to average the scores for different generation directions with different weights, which improves over the simple arithmetic averaging method used in BARTScore.
- 3) We proposed a novel entropy-based scheme for weighting the target generated terms, so that higher informative tokens receive more importance in accounting the score, which outperforms the naive uniform weighting employed in BARTScore.

2 Related work

2.1 Translation evaluation metrics

BLEU (Papineni et al., 2002) is the de facto metric for evaluating machine translation. It simply calculates n -gram matching between the reference and the hypothesis using precision scores with a brevity penalty. METEOR (Banerjee and Lavie, 2005) was developed to address two drawbacks of BLEU. It is F-score based (thus taking recall into account) and allows for a more relaxed matching, based on three forms: extract unigram, stemmed word, and synonym with WordNet (Miller, 1994). Apart from the above word-based metrics, some approaches operate at the character level. For example, chrF (Popović, 2015) computes the overall precision and recall over the character n -grams with various values of n . More recently, static word embeddings (Mikolov et al., 2013) have enabled capturing the semantic similarity between two texts possible, of what the historical metrics are incapable. Several metrics have been proposed to incorporate word vectors. For example, MEANT 2.0 (Lo, 2017) evaluates translation adequacy by measuring the similarity of the semantic frames and their role fillers between the human and machine translations.

Lately, pretrained language models have become popular, because they provide context-dependent

embeddings. This proved beneficial to all NLP tasks, but also to evaluation metrics. For example, using a modified version of the Word Mover’s Distance (Kusner et al., 2015), the Sentence Mover’s Similarity (Clark et al., 2019) measures the minimum cost of transforming one text into the other as the evaluation score, where sentences are represented as the average of their ELMo word embeddings (Peters et al., 2018). BERTR (Mathur et al., 2019) computes approximate recall based on the pairwise cosine similarity between the BERT word embeddings (Devlin et al., 2019) of two translations. UniTE (Wan et al., 2022) proposes a unified framework for modeling three evaluation prototypes: estimating the quality of the translation hypothesis by comparing it with reference-only, source-only, or source-reference-combined data. UniTE is built upon XLM-R multilingual language model (Conneau et al., 2020).

Among several alternatives, BERTScore (Zhang et al., 2019) and MoverScore (Zhao et al., 2019) have received more attention, and been adopted for reporting results in recent NLG publications (Lin et al., 2022; Weston et al., 2022). They both are unsupervised, general-purpose metrics and leverage BERT-like language models, however, with one difference lying in the similarity function for matching the two sequence representations. BERTScore greedily matches each token from one sequence to the single most similar token in the other sequence, in terms of cosine similarity of their token embeddings. While MoverScore conducts soft one-to-many matching using an n -gram generalization of the Word Mover’s Distance (Kusner et al., 2015).

Finally, the work closely related to ours is BARTScore (Yuan et al., 2021). Different with all the above metrics trying to match tokens or their embeddings, BARTScore proposes a novel conceptual view. It treats the evaluation of generated text as a text generation problem, with the help of a pretrained seq2seq model BART (Lewis et al., 2020). At the time of writing, this metric represents the state-of-the-art in the NLG evaluation. We will provide more details about it in Section 3.

2.2 Data augmentation

As deep learning models are often heavily reliant on large amounts of training data, a common attempt to get around the data scarcity problem is by applying data augmentation techniques (Shorten and Khoshgoftaar, 2019). These techniques in-

crease the size of the training set by making slightly modified copies of already-existing instances or by creating new, synthetic ones. Such augmented data have proven to be beneficial to the training of models in a wide variety of contexts, from computer vision (Shorten and Khoshgoftaar, 2019) to speech recognition (Bird et al., 2020), to NLP (Feng et al., 2021), as it acts as a regularizer and helps reduce overfitting (Krizhevsky et al., 2012). For dealing with textual data, a suite of augmentation techniques exists. To name only a few, backtranslation (Sennrich et al., 2016) translates a text into an intermediate language and then back into the original language, as a way of paraphrasing the initial text. Contextual augmentation (Kobayashi, 2018) generates augmented samples by randomly replacing words with others drawn following the in-context word distribution of a recurrent language model. SeqMix method (Guo et al., 2020) creates synthetic examples by softly mixing parts of two sentences via a convex combination.

Data augmentation has also been applied to the field of NLG evaluation metrics. BLEURT (Sellam et al., 2020) is a supervised metric, i.e., it requires to be finetuned on human meta-evaluations. Before finetuning, BLEURT creates an augmented synthetic dataset made by perturbing Wikipedia sentences with BERT mask-filling, backtranslation, and random word dropping techniques. The data are then annotated with some automatic numerical and categorical signals as pretraining labels. FrugalScore (Kamal Eddine et al., 2022) proposes the first knowledge distillation approach for NLG evaluation metrics, to alleviate the significant requirement of computational resources by the heavy metrics based on large pretrained language models (e.g., BERTScore and MoverScore). Unlike BLEURT, it is purely trained on a synthetic dataset consisting of pairs of more or less related sentences, created via various data augmentation techniques (e.g., paraphrasing with backtranslation, perturbation then denoising, etc.). The sentence pairs for training the student model are annotated with scores given by the metrics to be learned.

Differences. Note that BLEURT and FrugalScore use augmented data for the purpose of training their parameterized metric models, while our DATScore is an untrained and unsupervised metric not requiring human judgments for training and using augmented translation for the sole purpose of scoring.




Figure 1: Dashed arrows denote the generation directions covered by BARTScore. Solid black arrows indicate our newly introduced directions for calculating DATScore of the example *hypothesis* in English ($Hypo_{en}$). $Trans1_{xx}$ and $Trans2_{yy}$ represent data *augmented translations* in any languages xx and yy , obtained by applying a translation model (grey arrows) to the example *source* in French (Src_{fr}) and example *reference* in English (Ref_{en}), respectively.

3 DATScore

As mentioned in Subsection 2.1, BARTScore is not based on matching tokens nor their embeddings as the other evaluation metrics. Rather, it uses a novel approach by framing the evaluation of generated text as a text generation problem. Assuming first a pretrained seq2seq model is “perfect” (e.g., BART), BARTScore uses directly the model’s conditional probability of generating a provided target text Y given a provided input text X , as the evaluation score of the generation direction $X \rightarrow Y$. For example, Y corresponds to a translation hypothesis generated by any system, and X is the reference. If Y is of high quality, then by providing the pair to the pretrained BART model, the estimated conditional generation probability (evaluation score) $P(Y|X)$ should be high.

Therefore, with placing differently the *source* (Src), the *reference* (Ref), and the *hypothesis* ($Hypo$) in pair at the input or the output side of the trained seq2seq model for fetching conditional generation probabilities, BARTScore deviates three different generation directions illustrated as dashed arrows in Figure 1. The conditional probabilities associated with the directions are denoted as: Precision ($Ref \rightarrow Hypo$), Recall ($Hypo \rightarrow Ref$) and Faithfulness² ($Src \rightarrow Hypo$). Additionally, a F-score, the arithmetic average of Precision and Recall.

The score (conditional probability) for the generation direction from a source sequence $X =$

²BART being a monolingual model, faithfulness is only relevant in the context of abstractive summarization, and its corresponding direction cannot be applied to machine translation evaluation.

$\{x_t\}_{t=1}^n$ to a target sequence $Y = \{y_t\}_{t=1}^m$ is calculated as the factorized, weighted log probability over all generation steps:

$$\text{Score}_{X \rightarrow Y} = \sum_{t=1}^m w_t \log P(y_t | X, \{y_{t'}\}_{t'=1}^{t-1}; \theta) \quad (1)$$

where w_t denotes the term importance score to put different emphasis on different target tokens y_t . BARTScore simply employs a uniform weighting scheme (all equal to 1). θ denotes the parameterized seq2seq model.

Our contributions consist of three modifications tailored to machine translation:

Data augmented translations. Unlike BARTScore, we employ M2M-100 (Fan et al., 2021), a non-English-centric multilingual machine translation system as our backbone seq2seq model, due to its superior performance. As our main contribution, we translate the source (e.g., Src_{fr} in Figure 1) and the reference (Ref_{en}) into any languages as our augmented data (Trans1_{xx} and Trans2_{yy}) for evaluating the hypothesis (Hypo_{en}). In addition to the three directions covered by BARTScore, our metric takes into consideration all generation directions centered on the hypothesis connecting the source, the reference, and the two data augmented translations, i.e., in total 8 directions as the black (dashed and solid) arrows depicted in Figure 1. DATScore is calculated as the weighted average of the scores associated with all the directions:

$$\text{DATScore} = \sum_{X, Y} w_{X \rightarrow Y} \text{Score}_{X \rightarrow Y}; X \neq Y \quad (2)$$

where $w_{X \rightarrow Y}$ denotes the weight of the direction $X \rightarrow Y$, as detailed below.

One-vs-rest score averaging method. We observed empirically that sometimes, one direction score might strongly disagree with the others, likely being an outlier (failed evaluation). This may significantly affect the final DATScore correlations with the human meta-evaluations, if a simple arithmetic averaging method is applied (like BARTScore in computing F-score). To reduce this effect, we weight each direction with the sum of the Pearson correlations of its scores with the scores of all the other directions:

$$w_{X \rightarrow Y} = \sum_{X', Y'} \text{Corr}(\text{Score}_{X \rightarrow Y}, \text{Score}_{X' \rightarrow Y'}) \quad \text{s.t. } (X, Y) \neq (X', Y') \quad (3)$$

This one-vs-rest method will assign a low weight to the direction score correlated badly with the rest scores, thus reduce its negative effect to the averaging result.

Entropy-based term weighting scheme. BARTScore gives an equal weight w_t to every token in Equation 1 (uniform weighting). Instead, we introduce a novel scheme to give different importance to different target tokens y_t , based on the entropy:

$$w_t = - \sum_{i=1}^v P_t(z_i) \log P_t(z_i) \quad (4)$$

where v denotes the size of output generation vocabulary. $P_t(z_i)$ represents the probability of the i -th token in the vocabulary at time step t . Our assumption is that, when the model is very confident in generating the target token (low entropy), then this token is non-informative (e.g., stopword). While when the model is less confident (higher entropy), the target word is more informative, and then a higher weight should be assigned.

The effectiveness of all our choices regarding the above contributions is shown by our ablation studies (see Section 6).

4 Experiments

4.1 Experimental settings

We benchmark DATScore on two commonly used meta-evaluation datasets for machine translation metrics: WMT17 (Bojar et al., 2017) and WMT18 (Ma et al., 2018) consisting of multiple to_English and from_English language pairs. For each pair, a few thousand examples are available, each being made of a *source*, a *reference*, a *hypothesis* and a *label* produced by human annotators, assessing the quality of the system generated *hypothesis*. Depending on the *label* type, we use Kendall’s Tau τ correlations or absolute Pearson $|r|$ correlations. The former is used when relative ranking is provided, and the latter in the case of direct assessment. We adopt the Kendall’s Tau-like formulation proposed in (Bojar et al., 2017):

$$\tau = \frac{|\text{Concordant}| - |\text{Discordant}|}{|\text{Concordant}| + |\text{Discordant}|} \quad (5)$$

Metric	Model	$ r $:cs \rightarrow en	$ r $:de \rightarrow en	$ r $:fi \rightarrow en	$ r $:lv \rightarrow en	$ r $:ru \rightarrow en	$ r $:tr \rightarrow en	$ r $:zh \rightarrow en	Avg.
		/	/	/	/	/	/	/	
		τ :en \rightarrow cs	τ :en \rightarrow de	τ :en \rightarrow fi	τ :en \rightarrow lv	-	τ :en \rightarrow tr	-	
BLEU	1a) N/A	34.4/22.0	36.6/23.6	44.4/42.1	32.1/21.5	41.3/-	44.1/33.6	44.0/-	37.8/27.3
BERTScore	1b) RL/mBERT	71.0/43.8	74.5/40.4	83.3/58.8	75.6/46.6	74.6/-	75.1/57.1	77.5/-	75.9/49.3
MoverScore	1c) BB/mBERT	66.6/38.3	70.6/35.9	82.2/54.2	71.7/37.8	73.7/-	76.1/49.8	74.3/-	73.6/43.2
BARTScore	1d) BL+para/mBART	68.4/39.0	70.8/33.4	79.4/50.4	74.9/50.4	71.8/-	73.9/53.8	76.0/-	73.6/45.4
	1e) M2M-100_418M	65.9/45.0	66.1/44.5	79.9/59.2	71.7/40.3	69.0/-	71.8/70.9	71.6/-	70.9/52.0
	1f) M2M-100_1.2B	67.4/49.6	69.3/49.2	80.7/63.5	73.7/46.9	70.4/-	71.6/ 72.5	73.0/-	72.3/56.3
DATScore	1g) M2M-100_418M	68.6/51.1	68.5/48.1	82.0/63.7	74.7/48.3	73.0/-	77.6/70.9	76.5/-	74.4/56.4
	1h) M2M-100_1.2B	71.3/53.9	72.9/52.2	83.5/66.3	76.8/52.0	75.9/-	78.1/70.9	77.7/-	76.6/59.1

Table 1: Absolute Pearson correlation ($|r|$) for to-English and Kendall correlations (τ) for from-English with segment-level human scores on WMT17. BB stands of Bert-Base, RL for RoBERTa-Large and BL for BART-Large.

Metric	Model	τ :cs \rightarrow en	τ :de \rightarrow en	τ :et \rightarrow en	τ :fi \rightarrow en	τ :ru \rightarrow en	τ :tr \rightarrow en	τ :zh \rightarrow en	Avg.
		/	/	/	/	/	/	/	
		τ :en \rightarrow cs	τ :en \rightarrow de	τ :en \rightarrow et	τ :en \rightarrow fi	τ :en \rightarrow ru	τ :en \rightarrow tr	τ :en \rightarrow zh	
BLEU	2a) N/A	23.3/38.9	41.5/62.0	38.5/41.4	15.4/35.5	22.8/33.0	14.5/26.1	17.8/31.1	24.8/38.3
BERTScore	2b) RL/mBERT	40.4/55.9	55.0/72.7	39.7/58.4	29.6/53.9	35.3/42.4	29.2/38.9	26.4/36.1	36.5/51.2
MoverScore	2c) BB/mBERT	36.8/44.6	53.9/68.4	39.4/52.7	28.7/50.9	27.9/40.1	33.6/32.5	25.6/35.2	35.1/46.3
BARTScore	2d) BL+para/mBART	39.6/50.2	54.7/65.0	39.4/53.3	28.9/57.2	34.6/37.0	27.4/37.7	24.9/32.4	35.6/47.5
	2e) M2M-100_418M	36.3/55.4	53.5/72.2	37.6/58.4	26.3/60.2	33.4/44.4	26.8/45.1	23.4/31.3	33.9/52.4
	2f) M2M-100_1.2B	38.4/ 63.5	54.6/ 76.2	39.2/63.2	27.9/64.5	35.7/45.6	28.5/50.2	24.3/34.7	35.5/56.8
DATScore	2g) M2M-100_418M	38.6/53.5	53.5/71.3	39.3/64.0	28.4/62.2	34.9/44.4	28.5/47.9	25.3/34.0	35.5/53.9
	2h) M2M-100_1.2B	40.7/61.9	54.9/ 76.2	40.5/68.2	30.4/67.9	36.4/46.2	31.0/ 52.7	26.3/ 36.6	37.2/58.5

Table 2: Kendall correlations (τ) for to-English and from-English with segment-level human scores on WMT18. BB stands of Bert-Base, RL for RoBERTa-Large and BL for BART-Large.

where $|Concordant|$ is the number of examples on which the metric agrees with the human relative ranking, and $|Discordant|$ is the number of examples when they disagree.

To compute DATScore, two M2M-100 models: M2M-100_418M³ and M2M-100_1.2B⁴ are adopted (418M and 1.2B refer to the model sizes). They are finetuned to translate a source text to a target text by providing the source language code (e.g. "fr") at the beginning of the encoder input sequence, and a target language code at the beginning of the decoder input sequence. In our experiments, when English is the target language (to-English), we choose English for Trans1 and Spanish for Trans2 (see Figure 1). Otherwise, whenever English is the source language (from-English), we choose Spanish for Trans1 and English for Trans2. This choice is motivated by the fact that English and

Spanish are the top two represented languages in the training set of M2M-100 (Fan et al., 2021).

4.2 Main results

We compare the performance of our metric against BLEU and three other reference-based unsupervised metrics: BERTScore⁵, MoverScore⁶ and BARTScore⁷ (detailed in Subsection 2.1 and Section 3), using their official implementations. Experimental results are reported in Table 1 and 2. Following their original settings, we use different underlying language models for each baseline metric. For BERTScore and MoverScore, RoBERTa-Large (RL; Liu et al., 2019) and Bert-Base (BB) are used respectively when we evaluate to-English translations, and mBERT (Devlin et al., 2019) for from-English translations. In the case of BARTScore, we use a BART-Large (BL) checkpoint (finetuned on CNNDM (See et al., 2017) and

³https://huggingface.co/facebook/m2m100_418M

⁴https://huggingface.co/facebook/m2m100_1.2B

⁵https://github.com/Tiiiger/bert_score

⁶<https://github.com/AIPHES/emnlp19-moverscore>

⁷<https://github.com/neulab/BARTScore>

ParaBank2 (Hu et al., 2019) datasets) for evaluating to-English translations, and an mBART-50 model (Escolano et al., 2021) for from-English translations.

Overall, results show that on average across all language pairs, DATScore significantly outperforms all 4 baseline metrics under their original model settings (rows 1a-1d and 2a-2d). Specifically, with respect to the best performed baseline BERTScore (row 1b and 2b), our metric provides a performance boost of 0.7 for to-English case and of 9.8 for from-English case on WMT17 dataset in Table 1, and achieves a gain of 0.7 and of 7.3 respectively on WMT18 dataset in Table 2. These averaging results demonstrate the superiority and applicability of DATScore in evaluating general machine translations of many languages. Moreover, it’s interesting to note that our improvement is much more significant in from-English case, which makes DATScore particularly well-suited to evaluate hypothesis translations in non-English languages, often with low resource. We hypothesize that this is due to the inconsistency of underlying language models. The baselines adopt a monolingual model for evaluating English, but a multilingual one for non-English languages. However, DATScore uses a single multilingual M2M-100 model for both cases. It is known that, in general, monolingual models outperforms multilingual competitors. Thus, it is reasonable that when comparing multilingual-based DATScore against monolingual baselines in the to-English case, DATScore achieves a smaller improvement than in the other from-English case, where the comparison is fairer (multilingual vs. multilingual).

By looking across specific language pairs and directions, we observe DATScore constantly performs better than 4 baseline metrics with a few exceptions, i.e., de \rightarrow en (-1.6) in Table 1, and de \rightarrow en (-0.1), tr \rightarrow en (-2.6), and zh \rightarrow en (-0.1) in Table 2. Despite these small drops in the performance, DATScore brings a larger margin of improvement in most cases, such as en \rightarrow tr up to 13.8 both on WMT17 and WMT18 datasets.

In the end, for the sake of having a complete comparison, we additionally evaluate BARTScore⁸ with M2M-100_418M and M2M-100_1.2B models (row 1e, 1f, 2e, and 2f) that are used as DATScore’s underlying models. Results show that, only in

⁸The official implementation of BARTScore is slightly modified to take into account the languages tokens when using a multilingual model.

Metric	Model	WebNLG		
		SEMA	GRAM	FLU
BLEU	N/A	45.5	36.0	34.9
BERTScore	RoBERTa-Large	56.1	60.8	54.8
MoverScore	BERT-Base	-9.9	-27.8	-20.6
BARTScore	BART-Large+para	71.9	61.3	57.4
	M2M-100_418M	64.9	62.8	56.0
	M2M-100_1.2B	66.1	63.9	57.2
DATScore	M2M-100_418M	69.9	62.9	57.2
	M2M-100_1.2B	70.4	63.7	57.9

Table 3: Pearson correlation results on WebNLG dataset.

Metric	Model	REALSumm	SummEval			
		COV	COH	CONS	FLU	REL
BLEU	N/A	37.9	11.8	6.3	7.7	18.6
BERTScore	RoRERTa-Large	41.2	33.9	10.5	15.0	35.9
MoverScore	BERT-Base	44.1	14.4	14.7	13.8	29.1
BARTScore	BART-Large+para	31.7	20.8	-3.5	6.7	22.2
	M2M-100_418M	30.1	14.8	-2.3	3.0	19.8
	M2M-100_1.2B	32.0	17.1	1.1	6.7	22.8
DATScore	M2M-100_418M	44.7	17.1	4.4	4.6	26.3
	M2M-100_1.2B	45.5	19.5	6.8	8.2	30.2

Table 4: Pearson correlation results on two summarization datasets: REALSumm and SummEval.

the from-English case, while they bring improvement compared to the vanilla BARTScore (row 1d and 2d), they are not able to yield as big of a gain as our metric, indicating that our achieved improvement is not solely due to the underlying language model, but also to taking additional generation directions into account, including those related to data augmented translations.

5 Other NLG tasks

In addition to machine translation, our main focus, we evaluate DATScore on other NLG tasks, including data-to-text generation, abstractive summarization, and image captioning. To work around with the different modalities of source inputs represented in these tasks (e.g., not able to create a data augmented translation with an image), we adapt DATScore to only consider 4 generation directions: $Pred \leftrightarrow Ref$ and $Pred \leftrightarrow Trans2$.

Data-to-text. Table 3 shows the performance of DATScore compared to the other baselines on the WebNLG data-to-text dataset (Shimorina et al., 2018), which contains 2000 descriptions of structured tables along with their corresponding references. In addition, human assessments covering three dimensions are provided (*semantics*, *gram-*

mer, and fluency). The results show that DATScore significantly outperforms all the other metrics in two settings (grammer and fluency) out of three, while being very competitive in the third setting (semantics). Surprisingly, BERTScore is largely behind DATScore, and MoverScore failed to correlate positively with human judgments in all dimensions.

Summarization. Table 4 shows the evaluation of the different metrics on two summarization meta-evaluation datasets: REALSumm (Bhandari et al., 2020) and SummEval (Fabbri et al., 2021). Both datasets contain a few thousand examples of system-generated summaries and their references. The generated summaries are annotated with *lightweight pyramids* (Shapira et al., 2019) method in the case of REALSumm, while the annotations in SummEval cover four dimensions: *coherence, consistency, fluency, and relevance*. On REALSumm, DATScore has the best performance compared to all the other baselines even when using its smaller version (M2M-100_418M). However, despite its higher correlations compared to BARTScore and MoverScore, DATScore fails to outperform BERTScore on the different dimensions of SummEval.

Image captioning. We consider Flickr8K (Hodosh et al., 2013) and PASCAL-50S (Vedantam et al., 2015), two image captioning datasets. The former is annotated with scores from 1 to 4 assessing the relevance of the captions, and the latter is annotated with relative ranking (i.e., given two descriptions which one is better). Table 5 shows that in this task, DATScore is competitive to BARTScore and BERTScore. Surprisingly, MoverScore outperforms significantly all the other metrics despite its poor performance on the other datasets.

Finally, although not being the top performing metric across all tasks, DATScore showed an overall stable and competitive performance. Each of the other metrics fails in evaluating generations at least in one of the tasks. For example, BERTScore and MoverScore have a poor performance on the WebNLG dataset. On the other hand, BARTScore fails to correlate well on REALSumm and SummEval, even it is finetuned on an abstractive summarization dataset. This finding suggests that, regardless of being initially designed for machine translation evaluation, DATScore can be safely used to evaluate NLG systems in other tasks for different evaluation dimensions.

Metric	Model	Flickr8K	PASCAL-50S
		RELE	RR
BLEU	N/A	13.8	8.1
BERTScore	RoBERTa-Large	46.1	33.8
MoverScore	BERT-Base	52.5	33.2
BARTScore	BART-Large+para	44.8	33.1
	M2M-100_418M	34.3	29.6
	M2M-100_1.2B	34.6	26.3
DATScore	M2M-100_418M	42.6	29.6
	M2M-100_1.2B	45.3	31.4

Table 5: Pearson correlation Results on two Image Captioning datasets: Flickr8K and PASCAL-50S.

Entropy-based weighting	One-vs-rest weighting	to_English	from_English
✓	✓	37.2	58.5
✓	✗	37.1	58.1
✗	✓	36.4	55.9
✗	✗	36.4	56.0

Table 6: The average Kendall correlation (to/from)-English when the entropy-based and one-vs-rest weighting are included or excluded. Experiments are conducted on WMT18.

6 Ablation study

To validate our different choices with regards to DATScore, we conducted ablation studies on:

- 1) the contributions of all 8 direction scores, results are illustrated in Figure 2.
- 2) the effectiveness of our *one-vs-rest* score averaging and *entropy-based* term weighting strategies (See Section 3), results are reported in Table 6.

Contributions of all direction scores. From Figure 2(a), we observe that none of the individual directions (horizontal bars) has a better correlation with human judgements than DATScore (dashed vertical lines), which confirms the importance of our ensemble approach. In Figure 2(b), we can see that all variants excluding one direction will lead, in almost all cases, to a drop in the performance, compared to the complete DATScore in which all directions are included. Besides, in the case of to-English translations, we can see that the drop in the performance is almost the same for all exclusions of direction. While for from-English translations, the largest drop in performance is observed when *Hypo*→*Trans2* and *Trans2*→*Hypo* are excluded. This finding highlights the important contribution of our augmented data, especially in the




Figure 2: (a): The horizontal bars represent the Kendall correlations of **each individual generation direction**. (b): The horizontal bar represents the Kendall correlation of **a variant of DATScore with excluding the single generation direction** of the line. Both in (a) and (b), the dashed vertical lines represent the Kendall correlation of the vanilla and **complete DATScore**. Correlation results of to-English (in green) and from-English (in red) cases are calculated w.r.p human judgments, and averaged over all languages pairs. Experiments are conducted on WMT18.

low resource language settings (from-English). In the end, we can see that excluding *Src*→*Hypo* or *Trans1*→*Hypo* directions can lead to a slightly better final score. We leave the investigation of the potential negative impact of the two directions to a future work.

One-vs-rest and entropy-based weighting strategies. Table 6 shows the performance of DATScore variants with respect to different combinations of applying or not our proposed weighting strategies. Note that, when *one-vs-rest* and *entropy-based* weightings are not applied, they are replaced with a simple uniform averaging approach (as used in BARTScore). A performance drop is observed when excluding one of the two weighting strategies, especially for the entropy-based method, whose inclusion leads to an improvement of 2.5 compared to the uniform weighting. This experiment confirms the positive impact of our proposed weighting methods and motivates future work to further investigate a more elaborated approach in this direction.

7 Conclusion

In this work, we proposed one of the first applications of data augmentation techniques to NLG evaluation. To obtain an evaluation score of the translation hypothesis, our developed metric DATScore additionally leverages newly translated copies aug-

mented from the source and reference texts. We also proposed two novel strategies for score averaging and term weighting to improve the original, naive score computing process of BARTScore, on the basis of which our work is built. Experimental results show that DATScore achieved a higher correlation with human meta-evaluations, in comparison with the other recent state-of-the-art metrics, especially for those less represented languages other than English. Moreover, ablation studies show the effectiveness of our newly proposed score computing approaches, and extended experiments showed an overall stable and competitive performance of DATScore on more NLG tasks.

Limitations

In this section, we list some limitations that are worth further investigation in future works:

1) DATScore requires generating additional data augmented translations to perform evaluation. This process might be time consuming depending on the adopted backbone seq2seq model, especially if the original text is long. Thus, the performance scalability can be investigated in future complementary experiments.

2) We made a choice of using English and Spanish to create data augmented translations, for the reason that they are the most two represented languages in training M2M-100 model (see Subsec-

tion 4.1). This leaves a question to be answered about the performance of DATScore with augmentations varying in different other languages (e.g., Chinese). Moreover, for the sake of simplicity, we decided to only include a single translated copy of the source text and of the reference text each. However, this can be easily extended, and more augmented translations can be created in more languages. We expect to see an improvement in performance with diminishing returns.

3) BARTScore only considers the 8 generation directions centered on the hypothesis connecting with the source, the reference, and the two data augmented translations (see Section 3). However, other connections exist between these entities, such as $Src \rightarrow Ref$ and $Transl \rightarrow Src$ (see Figure 1). Therefore, future research could be dedicated to discovering the effect of these other directions and potentially leveraging them to improve the performance of DATScore.

4) Since our focus was on evaluating machine translation, we naturally chose translation for augmenting the data. However, other data augmentation techniques could be seamlessly integrated into DATScore, such as using a text paraphrasing model (Bandel et al., 2022).

References

- Elron Bandel, Ranit Aharonov, Michal Shmueli-Scheuer, Ilya Shnayderman, Noam Slonim, and Liat Ein-Dor. 2022. [Quality controlled paraphrase generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 596–609, Dublin, Ireland. Association for Computational Linguistics.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Manik Bhandari, Pranav Narayan Gour, Atabak Ashfaq, Pengfei Liu, and Graham Neubig. 2020. [Re-evaluating evaluation in text summarization](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9347–9359, Online. Association for Computational Linguistics.
- Jordan J Bird, Diego R Faria, Cristiano Premevida, Anikó Ekárt, and Pedro PS Ayrosa. 2020. Overcoming data scarcity in speaker identification: Dataset augmentation with synthetic mfccs via character-level rnn. In *2020 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, pages 146–151. IEEE.
- Ondřej Bojar, Yvette Graham, and Amir Kamran. 2017. [Results of the WMT17 metrics shared task](#). In *Proceedings of the Second Conference on Machine Translation*, pages 489–513, Copenhagen, Denmark. Association for Computational Linguistics.
- Elizabeth Clark, Asli Celikyilmaz, and Noah A. Smith. 2019. [Sentence mover’s similarity: Automatic evaluation for multi-sentence texts](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2748–2760, Florence, Italy. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- 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.
- Carlos Escolano, Marta R. Costa-jussà, José A. R. Fonollosa, and Mikel Artetxe. 2021. [Multilingual machine translation: Closing the gap between shared and language-specific encoder-decoders](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 944–948, Online. Association for Computational Linguistics.
- Alexander R. Fabbri, Wojciech Kryściński, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir Radev. 2021. [SummEval: Re-evaluating summarization evaluation](#). *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, et al. 2021. Beyond english-centric multilingual machine translation. *J. Mach. Learn. Res.*, 22(107):1–48.
- Steven Y. Feng, Varun Gangal, Jason Wei, Sarath Chandar, Soroush Vosoughi, Teruko Mitamura, and Edouard Hovy. 2021. [A survey of data augmentation approaches for NLP](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*,

- pages 968–988, Online. Association for Computational Linguistics.
- Demi Guo, Yoon Kim, and Alexander Rush. 2020. [Sequence-level mixed sample data augmentation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5547–5552, Online. Association for Computational Linguistics.
- Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899.
- J. Edward Hu, Abhinav Singh, Nils Holzenberger, Matt Post, and Benjamin Van Durme. 2019. [Large-scale, diverse, paraphrastic bitexts via sampling and clustering](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 44–54, Hong Kong, China. Association for Computational Linguistics.
- Moussa Kamal Eddine, Guokan Shang, Antoine Tixier, and Michalis Vazirgiannis. 2022. [FrugalScore: Learning cheaper, lighter and faster evaluation metrics for automatic text generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1305–1318, Dublin, Ireland. Association for Computational Linguistics.
- Sosuke Kobayashi. 2018. [Contextual augmentation: Data augmentation by words with paradigmatic relations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 452–457, New Orleans, Louisiana. Association for Computational Linguistics.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *International conference on machine learning*, pages 957–966. PMLR.
- 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.
- 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.
- Haitao Lin, Junnan Zhu, Lu Xiang, Yu Zhou, Jiajun Zhang, and Chengqing Zong. 2022. [Other roles matter! enhancing role-oriented dialogue summarization via role interactions](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2545–2558, Dublin, Ireland. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Chi-kiu Lo. 2017. [MEANT 2.0: Accurate semantic MT evaluation for any output language](#). In *Proceedings of the Second Conference on Machine Translation*, pages 589–597, Copenhagen, Denmark. Association for Computational Linguistics.
- Qingsong Ma, Ondřej Bojar, and Yvette Graham. 2018. [Results of the WMT18 metrics shared task: Both characters and embeddings achieve good performance](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 671–688, Belgium, Brussels. Association for Computational Linguistics.
- Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2019. [Putting evaluation in context: Contextual embeddings improve machine translation evaluation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2799–2808, Florence, Italy. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- George A. Miller. 1994. [WordNet: A lexical database for English](#). In *Human Language Technology: Proceedings of a Workshop held at Plainsboro, New Jersey, March 8-11, 1994*.
- Bonan Min, Hayley Ross, Elior Sulem, Amir Pouran Ben Veyseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heinz, and Dan Roth. 2021. Recent advances in natural language processing via large pre-trained language models: A survey. *arXiv preprint arXiv:2111.01243*.
- Jun-Ping Ng and Viktoria Abrecht. 2015. [Better summarization evaluation with word embeddings for ROUGE](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1925–1930, Lisbon, Portugal. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia,

- Pennsylvania, USA. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Maxime Peyrard. 2019. [Studying summarization evaluation metrics in the appropriate scoring range](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5093–5100, 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.
- Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. 2020. Pre-trained models for natural language processing: A survey. *Science China Technological Sciences*, 63(10):1872–1897.
- Ananya B Sai, Akash Kumar Mohankumar, and Mitesh M Khapra. 2022. A survey of evaluation metrics used for nlg systems. *ACM Computing Surveys (CSUR)*, 55(2):1–39.
- 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. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Ori Shapira, David Gabay, Yang Gao, Hadar Ronen, Ramakanth Pasunuru, Mohit Bansal, Yael Amsterdamer, and Ido Dagan. 2019. [Crowdsourcing lightweight pyramids for manual summary evaluation](#). 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 682–687, Minneapolis, Minnesota. Association for Computational Linguistics.
- Anastasia Shimorina, Claire Gardent, Shashi Narayan, and Laura Perez-Beltrachini. 2018. *WebNLG challenge: Human evaluation results*. Ph.D. thesis, Loria & Inria Grand Est.
- Connor Shorten and Taghi M Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Yu Wan, Dayiheng Liu, Baosong Yang, Haibo Zhang, Boxing Chen, Derek Wong, and Lidia Chao. 2022. [UniTE: Unified translation evaluation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8117–8127, Dublin, Ireland. Association for Computational Linguistics.
- Jack Weston, Raphael Lenain, Udeepa Meepegama, and Emil Fristed. 2022. [Generative pretraining for paraphrase evaluation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4052–4073, Dublin, Ireland. Association for Computational Linguistics.
- Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: Evaluating generated text as text generation. *Advances in Neural Information Processing Systems*, 34:27263–27277.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. 2019. [MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance](#). 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 563–578, Hong Kong, China. Association for Computational Linguistics.