

Perturbation CheckLists for Evaluating NLG Evaluation Metrics

Ananya B. Sai^{1,2} Tanay Dixit¹ Dev Yashpal Sheth^{1,2}
Sreyas Mohan³ Mitesh M. Khapra^{1,2,4}

¹ Indian Institute of Technology, Madras

² Robert Bosch Centre for Data Science and Artificial Intelligence, IIT Madras

³ Center for Data Science, New York University ⁴ AI4Bharat

{ananya, devsheth, miteshk}@cse.iitm.ac.in,

dixittanay@gmail.com , sm7582@nyu.edu

Abstract

Natural Language Generation (NLG) evaluation is a multifaceted task requiring assessment of multiple desirable criteria, *e.g.*, fluency, coherency, coverage, relevance, adequacy, overall quality, *etc.* Across existing datasets for 6 NLG tasks, we observe that the human evaluation scores on these multiple criteria are often not correlated. For example, there is a very low correlation between human scores on *fluency* and *data coverage* for the task of structured data to text generation. This suggests that the current recipe of proposing new automatic evaluation metrics for NLG by showing that they correlate well with scores assigned by humans for a *single* criteria (overall quality) alone is inadequate. Indeed, our extensive study involving 25 automatic evaluation metrics across 6 different tasks and 18 different evaluation criteria shows that there is no single metric which correlates well with human scores on *all* desirable criteria, for most NLG tasks. Given this situation, we propose *CheckLists* for better design and evaluation of automatic metrics. We design templates which target a specific criteria (*e.g.*, coverage) and perturb the output such that the quality gets affected only along this specific criteria (*e.g.*, the coverage drops). We show that existing evaluation metrics are not robust against even such simple perturbations and disagree with scores assigned by humans to the perturbed output. The proposed templates thus allow for a fine-grained assessment of automatic evaluation metrics exposing their limitations and will facilitate better design, analysis and evaluation of such metrics.¹

1 Introduction

As the number of tasks and benchmarks for NLG have increased (Gehrmann et al., 2021), the challenges in evaluating NLG systems have also continued to grow (Liu et al., 2016; Nema and Khapra,

2018; Sai et al., 2019). One reliable way of evaluating NLG systems is to collect human judgements. However, this is a time consuming and expensive process (Freitag et al., 2021; Deriu et al., 2019; Howcroft et al., 2020). Hence, automatic evaluation metrics such as BLEU (Papineni et al., 2002) which are quicker to compute have become popular, despite being less reliable (Callison-Burch et al., 2006; Reiter, 2018).

The survey by Sai et al. (2020b) shows that more than 35 automatic evaluation metrics have been proposed for NLG since 2014, however, there is no careful evaluation of the ability of such metrics to assess the quality of the output of an NLG system on multiple desired criteria. For example, consider the task of dialog evaluation, where humans are asked to score the output on multiple criteria such as *fluency*, *adequacy*, *coherence*, *informativeness*, *engagingness*, *consistency*, *etc.* Contrast this with automatic evaluation metrics such as BLEU, BLEURT (Sellam et al., 2020), DEB (Sai et al., 2020a), ADEM (Lowe et al., 2017), *etc.*, which assign a *single* score to the output. What does this score indicate? More specifically, does a low DEB score indicate that the output is not fluent or does it indicate that the output is fluent but not coherent or neither fluent nor coherent? Hence, a single overall score assigned by automatic evaluation metrics is not very informative in deciding which aspects of model improvement should one focus on.

The question then is why do current automatic evaluation metrics produce only a single overall score? This is simply because of a conscious choice made while designing automatic evaluation metrics. In particular, current works only focus on evaluating whether the scores assigned by the proposed metric correlate well with the *overall quality* scores assigned by humans as opposed to *all* relevant criteria. In this work, we make a case for shifting the focus to *all* relevant criteria while evaluating such metrics. To this end, we first do a systematic study

¹Our templates and code are available at <https://iitmnlp.github.io/EvalEval/>

Original sentence:	Cameron is the director of Titanic
Change names:	Kate is the director of Titanic (incorrect)

Table 1: Example of a perturbed output.

involving 6 NLG tasks, 18 different human evaluation criteria (*fluency, coverage, coherence, consistency, etc*) and 25 automatic evaluation metrics. We take existing English datasets containing human judgements for various tasks and criteria and make two important observations. First, for a given task, human scores for different criteria often have a low correlation, thereby suggesting that these criteria cannot be clubbed together and evaluated using a single score assigned by an automatic evaluation metric. Second, none of the automatic evaluation metrics have a high correlation with human scores for any of the desired criteria for a given task.

The above results highlight a lacunae in the evaluation of automatic evaluation metrics wherein their ability to assess the output on multiple criteria is not evaluated. In this work, we propose a flexible framework which allows a systematic evaluation of the capabilities of an automatic evaluation metric. In particular, we propose CheckList-style templates (Ribeiro et al., 2020) which evaluate the robustness of the metrics to certain perturbations targeting specific criteria. We illustrate this idea with an example in Table 1. In row 2 of Table 1 a gold standard output is perturbed by changing named entities, thereby affecting its *factual correctness* which is important for data-to-text generation. If an automatic evaluation metric indeed evaluates *factual correctness* then its score should drop when presented with such a perturbed output.

For the 6 NLG tasks mentioned earlier, we create 34 such perturbation templates covering 18 different evaluation criteria. We then instantiate these templates to create large-scale test cases. For every perturbation, we also collect human judgements to understand how much would a human change his/her score when shown such a perturbed output. We find that for several perturbations, the scores assigned by automatic evaluation metrics do not agree with the scores assigned by humans, thereby indicating that current automatic evaluation metrics are not robust to such perturbations (i.e., they do not really evaluate the desired criteria). Overall, we believe that the proposed templates provide a better framework for a more fine-grained evaluation of automatic evaluation metrics which goes much beyond computing correlations with human scores.

2 Criteria used in Human Evaluations

The goal of this work is to carefully evaluate automatic evaluation metrics with a focus on their ability to capture the diverse set of criteria used by humans while assessing NLG systems. To begin with, we describe the criteria that the output of an NLG system must satisfy for the 6 NLG tasks that we consider in this work, viz., machine translation (MT), dialog generation (DG), automatic summarisation (AS), question generation (QG), data-to-text generation (D2T) and image captioning (IC). Over the years, different works have proposed different criteria for evaluating NLG systems. In this work, we consider a popular set of criteria for each task as summarised in Sai et al. (2020b) and presented in Table 2. Given the wide variety of criteria used for each task, one obvious question to ask is whether we really need so many criteria or is a single overall score enough. One could argue that it is obvious from the definitions of the criteria that each of them is unique and a good score on one (say, fluency) may not necessarily imply a good score on another (say, coverage). However, we provide a quantitative argument for this by computing the correlations between human scores for different criteria as described below.

2.1 Correlations between different criteria

We use existing publicly available datasets containing human judgement scores on multiple criteria for each of the 6 tasks described earlier. For example, (Castro Ferreira et al., 2020) contains 3025 samples of outputs generated by data-to-text generation systems that participated in the WebNLG 2020 challenge. For each of these samples, the organisers asked humans to rate the output based on 5 criteria, viz., *fluency, data coverage, relevance, correctness* and *text structure*. We use these scores to compute the correlations between the scores of all the $\binom{5}{2}$ pairs of criteria. We repeat this for the other tasks using the datasets described in Table 3^{2 3} Using these annotations, we compute the pairwise Kendall tau correlations between all criteria

²For AS, we could not find a dataset containing human judgements for the set of criteria in Sai et al. (2020b). Hence, we use the dataset provided by Fabbri et al. (2020).

³Note that all of the datasets mentioned in Table 3 were collected using well established methods to ensure that the annotations were of high quality. Some of these datasets do not explicitly report the Inter Annotator Agreement (IAA) scores whereas others (Fabbri et al., 2020; Castro Ferreira et al., 2020; Nema and Khapra, 2018) report a good IAA score ranging from 0.63-0.71.

Task	Criteria
Machine Translation	Adequacy: The generated translation should adequately represent all the information present in the reference.
Question Generation	Relevance: Is the question related to the source material they are based upon. Answerability: Is the generated question answerable given the context.
Abstractive Summarization	Informativeness: The summary should convey the key points of the text. Non-redundancy: The summary should not repeat any points, and ideally have maximal information coverage within the limited text length. Referential clarity: Any intra-sentence or cross-sentence references in the summary should be unambiguous and within the scope of the summary. Focus: The summary needs to have a focus and all the sentences need to contain information related to this focal point. Structure and Coherence: The summary should be a well-organized and coherent body of information
Dialogue Generation	Making sense: Does the bot say things that don't make sense? Engagingness: Is the dialogue agent enjoyable to talk to? Interestingness: Did you find the bot interesting to talk to? Inquisitiveness: Does the bot ask a good amount of questions? Listening: Does the bot pay attention to what you say? Avoiding Repetition: Does the bot repeat itself? (either within or across utterances) Humanness: Is the conversation with a person or a bot?
Image Captioning	Relevance: The caption should be specific and related to the image. Thoroughness: The caption should adequately describe the image.
Data to Text Generation	Data Coverage: Does the text include descriptions of all predicates presented in the data? Relevance: Does the text describe only such predicates which are found in the data? Correctness: When describing predicates which are found in the data, does the text mention correct the objects and adequately introduces the subject for this specific predicate? Text Structure: Is the text grammatical, well-structured, written in acceptable English?
All above tasks	Fluency: How fluent is the generated text?

Table 2: Criteria used for human judgements to evaluate NLG systems for each of the 6 tasks

Task	Data collected/ released by	# Samples (annotators)
QG	Nema and Khapra (2018)	1000 (in-house)
AS	Fabbri et al. (2020)	1600 (expert, crowdsourced)
D2T	Castro Ferreira et al. (2020)	3025 (crowdsourced)
DG	See et al. (2019)	3316 (crowdsourced)
MT	Callison-Burch et al. (2007)	10,754 (crowdsourced)
IC	Aditya et al. (2015) (Coco subset)	2007 (crowdsourced)

Table 3: Datasets with human scores on many criteria.

for each task as seen in Figure 1. (Refer to appendix A for pearson correlations of the criteria.)

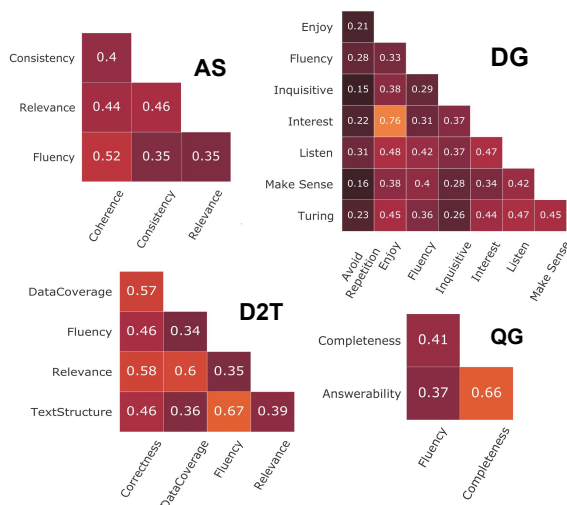


Figure 1: Correlations between criteria for different tasks. **The darker the shade of the cell, the lower is the correlation.**

We see that, across tasks, for most pairs of criteria, the correlation is moderate (between 0.3 and 0.5) to low (< 0.3). The highest correlation of 0.76 is observed between *interestingness* and *enjoyability* for dialogue generation. However other criteria such as avoiding *repetition*, *inquisitiveness*, and *making sense* have low correlations with most of the other criteria. We make similar observations for the correlations between the criteria for other tasks. Even for IC the correlation between the 2 criteria of thoroughness and correctness is 0.41. For MT, the commonly used criteria of *fluency* and *adequacy* were found to be highly correlated with Pearson correlation co-efficient of 0.69 (Banchs et al., 2015). This justifies why WMT evaluations now ask humans to give to only a single score indicating overall quality. However, given the low to moderate correlations between criteria for other tasks, a similar strategy is not prudent for these tasks.

Takeaway: For tasks whose linguistic criteria show a low or at-best moderate correlation with each other, a single score assigned by an automatic metric is inadequate for a comprehensive assessment.

3 Perturbation Checklists

So far we have established that if automatic evaluation metrics are to be used as a substitute for human evaluations as a whole, then they should be capable

Task	Criteria	Perturbation	Unmodified sentence	Perturbed sentence	
All tasks	Fluency	Jumbling word order	We play badminton every evening.	We badminton every evening play.	
		Subject-verb disagreement	He doesn't know how to bake.	He don't know how to bake.	
		Spelling errors	Make the most of every opportunity presented to you.	Make the most of evry opportunity presented to you.	
	Invariance	Replace with synonyms	The mangoes are delicious .	The mangoes are tasty .	
		Contractions	We are going to embark on an adventure.	We're going to embark on an adventure.	
MT	Adequacy	Numerals to words	The flight will be delayed by 2 hours.	The flight will be delayed by two hours.	
		Dropping out words or phrases	I was being followed.	I followed.	
		Add extra text	This book is so inspiring.	This book is so inspiring, I forgot .	
AS	Informativeness	Negation/antonyms	It will rain on Monday.	It will not rain on Monday.	
		Use hyponyms to create misinformation	The girl my brother Andy met through MySpace turned out to be completely made up.	The girl my friend Andy met through MySpace turned out to be completely made up.	
	Flow / coherence	Reorder sentences	The pandemic was spreading uncontrollably. Vaccines are being developed and tested rapidly.	Vaccines are being developed and tested rapidly. The pandemic was spreading uncontrollably.	
QG	Answerability	Referential clarity	Replace nouns by pronouns	The pandemic was spreading uncontrollably. Vaccines are being developed rapidly.	It was spreading uncontrollably. They are being developed rapidly.
		Drop question word	When was he born ?	Was he born?	
	Change question to assertive statement	Who is the director of Titanic?	The director of Titanic is James Cameron.		
DG	Relevance	Mask a few words and predict	How could Tesla run the experiments ?	How could Tesla run to the beach ?	
		Perturb nouns	Why did Mary go to the school ?	Why did Raj go to the market ?	
	Making sense	Negate a previous statement by same speaker	Bot: I enjoy having your daughter in my class. User: I'm glad to hear that.	I don't enjoy having your daughter in my class.	
IC	Avoid repetition	Add extra non-sensible text	Do you know where Dr. XYZ lives?	Yes, my father is my grandmother's father	
		Repeat previous utterances	Do you know Dr.XYZ?	Do you know where Dr. XYZ lives?	
	Repeat phrases	I like ice creams	I like ice creams, ice creams		
D2T	Data Coverage	Listening	Replace with "Can you repeat?"	I need to book a taxi	I'm sorry, can you repeat?
		Relevance	Map random responses	I am new to coding .	I am scared of snakes .
	Correctness	Change gender	Change attributes	Two girls are playing with a doll	Two boys are playing with a doll
A small boy playing with a red ball			A tall boy playing with a green ball		
Drop objects/noun		A small boy playing with a red ball	A small playing with a red		
Thoroughness	Repeat (append) object	A lady riding a horse.	A lady riding a horse and a lady .		
		Use hyponyms	Beethoven was a German musician	Beethoven was a German architect	
	Change numbers	The cricketer was born in 1990 .	The cricketer was born in 1950 .		
Relevance	Perturb names	Drop phrases	A small boy playing with a red ball	A small boy playing with a	
		Repeat phrases	Beethoven was a German musician	Beethoven was a German musician and German musician .	

Table 4: Perturbation templates targeting various criteria with examples. The blue highlights indicate the portions of the original sentence affected by the perturbation template. The red highlights indicate the changes in the modified sentence.

of evaluating the output on multiple desired criteria. However, the current recipe of proposing and evaluating evaluation metrics does not take this into account. To enable such a systematic evaluation of automatic evaluation metrics, we propose *perturbation checklists*. Similar to the original Checklist paper (Ribeiro et al., 2020), the idea is to evaluate the performance of the evaluation metric in detecting criteria-specific changes in the output.

We design such perturbation templates for each relevant criteria for each of the 6 tasks as shown in Table 4. For example, consider the criteria *fluency* which is relevant for all the tasks. Now consider a perturbation template for this criteria which simply drops the stop words in the output. Such a perturbation would definitely affect the fluency of the output. If an automatic evaluation metric is capable of assessing fluency, then this drop in the fluency of the output should get reflected in the score assigned by the metric. More formally, let p be the original output and \hat{p}_c^t be the output obtained by applying the perturbation template t for the criteria c . Further, let $f_e(p)$ be the score assigned by a given

evaluation metric e to the output p , normalised to be in the range $[0, 1]$. If the metric e is capable of assessing fluency then we would expect $f_e(\hat{p}_c^t)$ to be lower than $f_e(p)$. Now, further let $h(p)$ and $h(\hat{p}_c^t)$ be the scores (also normalised to have range $[0, 1]$) assigned to the original and perturbed outputs by human annotators. We then define a metric $s_c^t(e)$ which captures the ability of the metric e to detect the perturbation t for the desired criteria c .

$$s_c^t(e) = (h(\hat{p}_c^t) - h(p)) - (f_e(\hat{p}_c^t) - f_e(p)) \quad (1)$$

The score $s_c^t(e)$ as defined above thus captures the deviation between a human's perception about the drop in the quality and the metric e 's perception about the drop in the quality.

We design a total of 34 such perturbation templates across all the criteria and all the tasks. Each template is manually created by us and targets a specific criteria. We also present *invariant* templates that do not affect any criteria although they modify the sentences. For perturbations resulting from such invariant templates the score of the metric should not drop. The invariant and fluency-based

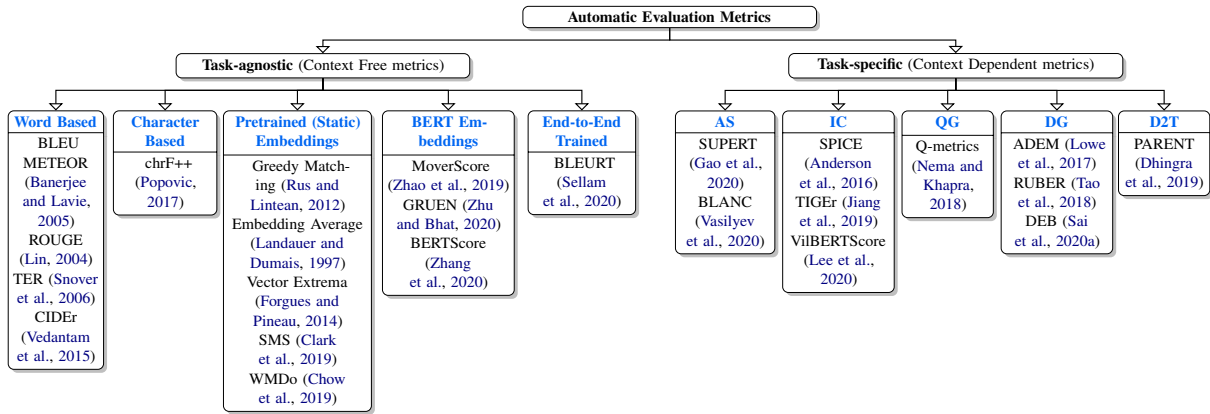


Figure 2: Metrics analysed in this study

templates are common for all the tasks considered in this work. Table 4 shows sample perturbations generated by each of the templates. (Please refer to appendix C for a more comprehensive list of the proposed perturbations with examples for each task). These perturbed sentences are generated automatically using the checklist framework (Ribeiro et al., 2020). This framework contains modules for performing simple string manipulations such as dropping stop words, replacing/dropping named entities, masking words or replacing them by other words/phrases. We also extend the framework with additional modules for jumbling words, changing numbers to words, subject-verb disagreement, changing gender, reordering sentences, adding spurious text, and adding redundancy at the word/phrase/sentence-level.

4 Experimental setup

We first do a coarse grained evaluation of several metrics by computing their correlations with the scores assigned by humans for multiple criteria. Note that unlike existing studies which study such correlations for a small number of metrics (typically, n -gram based metrics) for a specific task (say, MT) and a single criteria (typically, overall quality), we do a more comprehensive study involving a combination of 6 tasks, 25 metrics and multiple criteria. Apart from this coarse grained evaluation which simply looks at correlations, we also do a more fine-grained evaluation of the robustness of these metrics to different criteria-specific perturbations as summarised in Table 4. This fine-grained evaluation augments the coarse-grained evaluation and helps us understand the evaluation capabilities of these metrics. Below, we describe the datasets and evaluation metrics used in our work.

Datasets. For the coarse-grained evaluation, we use the datasets containing human judgements as described earlier in Table 3 in Section 2. For the fine-grained evaluation, we use datasets containing multiple ground truth references which can then be perturbed using our templates. For MT, we use the expanded version of newstest2017 Chinese to English dataset (Hassan et al., 2018) which contains two references for each sentence. For QG, we use the SQuAD dataset (Rajpurkar et al., 2016) which contains multiple questions for each passage. For AS, we use the curated personal narrative corpus (Ouyang et al., 2017). For DG, we use DailyDialog++ (Sai et al., 2020a) which contains two-speaker conversations on generic topics. For IC, we use the COCO component of the Composite dataset containing 5 reference captions for each image (Aditya et al., 2015). Lastly, for D2T, we use the Triples-to-Text data of the WebNLG 2020 challenge dataset (Castro Ferreira et al., 2020).

Applying perturbations. We take the reference sentences from the above task-specific datasets and apply perturbations using the Checklist framework described earlier in section 3. We first preprocess the sentences by performing tokenization, part-of-speech tagging, named entity recognition, *etc.* The targeted part of the sentence is then modified either by leveraging simple string manipulation functions or by masking and generating the words/phrases using the predictions by RoBERTa (Liu et al., 2019). We provide more implementation details in appendix B.

Automatic Evaluation Metrics. We study a total of 25 evaluation metrics belonging to different classes as shown in Figure 2. For BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), ROUGE-L (Lin, 2004), CIDEr (Vedantam

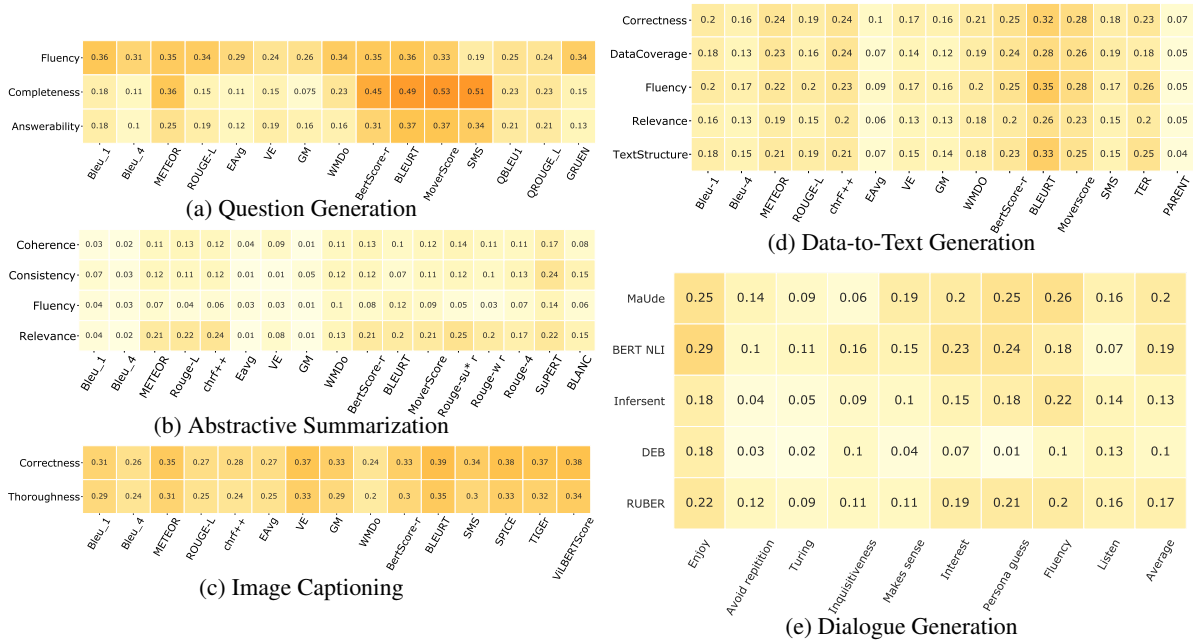


Figure 3: Correlations of metrics with different criteria (Note: For MT, we refer to the results of the WMT 2020 shared metrics task (Mathur et al., 2020b). For DG, we only present the reference-less metrics since the dataset does not contain references to compute the task-agnostic metric scores.) (Refer appendix D for full plots)

et al., 2015), Embedding Averaging (EAvG) (Laudauer and Dumais, 1997), Greedy Matching (GM) (Rus and Lintean, 2012), and Vector Extrema (VE) (Forgues and Pineau, 2014), we use the implementation provided by Sharma et al. (2017). For chrF++ (Popovic, 2017), TER (Snover et al., 2006), BERTScore (Zhang et al., 2020), and BLEURT (Sellam et al., 2020) we use the repository of Castro Ferreira et al. (2020). For SMS (Clark et al., 2019), WMD (Chow et al., 2019), and MoverScore (Zhao et al., 2019), we use the implementation provided by Fabbri et al. (2020). For all the task-specific metrics in Figure 2, we use the official codes from the respective papers.

Collecting human judgements. For eq 1, we need human judgement scores. We collect these with the help of 15 annotators who were computer science graduates with a background in the field of Natural Language Processing (NLP) and are also proficient in English language. For each task and criteria, the annotators were provided with the corresponding perturbation templates and asked to provide a penalty score indicating by how much a perturbation would alter the meaning/essence of a sentence on a scale of 0-10. A score of 0 indicates that there is no difference between the original and modified sentences upon application of the perturbation whereas a score of 10 indicates that the perturbation drastically alters sentences. These scores are nor-

malised to be in the range $[0, 1]$. The term $h(\hat{p}_c^t)$ in eq 1 for perturbation t is computed by subtracting the mean of all normalised human penalty scores from 1.

The standard deviation of the normalised annotator scores for each perturbation lies between 0.03 to 0.2. We also measure the inter annotator agreement by splitting annotators randomly into 2 groups and computing the kendall tau correlation score between the average scores of the 2 groups following Liu et al. (2016). This process was run 5 times with different seeds and the final inter-annotator correlation score was found to be 0.79.

5 Results and Discussion

We now discuss the results of our experiments.

5.1 Insights from coarse-grained evaluation

Figure 3 shows the correlations of different automatic evaluation metrics with multiple evaluation criteria, across 6 different NLG tasks. Our main observations from the figure are:

Mostly poor correlations of metrics across criteria. We observed that across all tasks and all criteria, most of the metrics have poor correlations. In particular, out of the 271 correlation values reported in Figure 3, 228 are poor (<0.3), 35 are moderate (between 0.3 and 0.5) and only 2 are high (>0.5). Surprisingly even some very recently

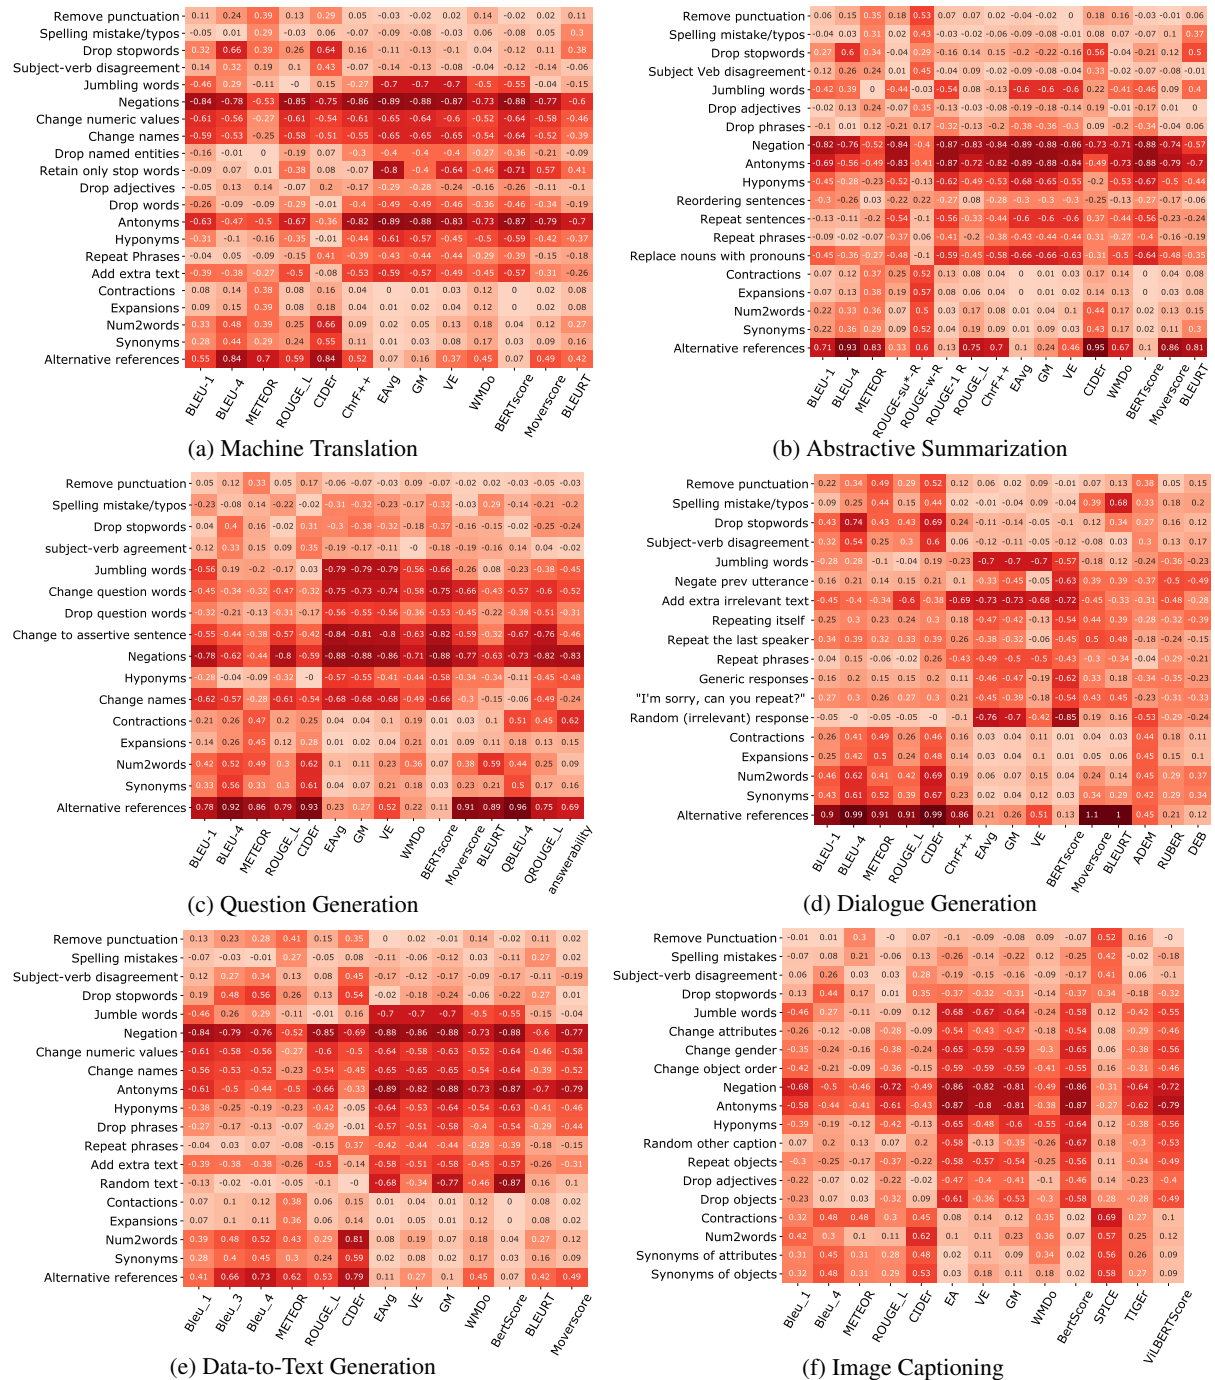


Figure 4: Heatmap of deviation of metric scores from human averages. The darker the cell, the more the absolute deviation from human scores.

proposed metrics such as DEB (Sai et al., 2020a), BLANC (Vasilyev et al., 2020) and MaUde (Sinha et al., 2020) do not correlate well with human judgements on other criteria. This is despite the fact that these are task-specific metrics which use the modern machinery of pre-trained BERT-based models and are fine-tuned on human judgements for overall quality. This vindicates our stand that simply tuning for overall quality does not lead to good

correlations with other criteria. We do observe a few decently correlated metrics for some of the tasks along a few criteria. Specifically, the moderate correlations are found (i) in D2T for BLEURT along the dimensions of fluency, correctness and text structure, (ii) in IC for all of the task-specific metrics, majority of the embedding-based metrics and for METEOR, (iii) in QG for most metrics (except SMS) along fluency, for BERTScore, Mover-

score, SMS and BLEURT along answerability, as well as along completeness, on which we find some of the highest correlations.

Pre-training and/or training often helps. While almost all the metrics have poor correlations with different criteria across tasks, we observe that the ones which use a pre-trained component such as static or contextualised word/sentence embeddings and/or use task-specific training data perform better. For example, BLEURT which uses a pre-trained BERT and is fine-tuned using human judgements for MT, is among the top performing metrics across all the tasks. These findings are also consistent with those reported in the WMT20 shared task on evaluation metrics for MT (Mathur et al., 2020b).

Task-agnostic metrics versus task-specific metrics. For the tasks of QG, AS, IC and D2T we find that task-agnostic metrics such as BERTScore, Moverscore and BLEURT are consistently among the top performing metrics (i.e., their correlation scores are either the best or close to the best scores for a given task and criteria) along with the task-specific metrics. This is interesting as these task-agnostic metrics were not fine-tuned on any task specific human judgements and were originally proposed for a different task (MT). Among the task-specific metrics, SUPERT for AS has a relatively better correlation with consistency than all other metrics. Overall, there seems to be scope for more work/improvements on task-specific metrics to capture the criteria peculiar to each task.

5.2 Insights from fine-grained evaluation

We now complement the above analysis with a more fine-grained analysis using perturbation templates. To do so, we use the perturbation templates described in Section 3 and plot the deviation between metric scores and human scores using the formula in Equation 1. These results are presented in Figure 4 and summarised below.

Correlations do not reveal everything. In the previous section we observed that BERT-based metrics such as BERTScore, BLEURT and MoverScore are among the top performing metrics across tasks and criteria. However, our analysis with perturbation templates reveals that even these metrics are not robust to very simple perturbations. For example, for the task of MT consider the perturbations of adding negations, changing names, changing numeric values or replacing by antonyms in the output which can significantly alter the meaning of

the sentence and thereby affect adequacy. However, BERTScore, BLEURT and MoverScore are not able to perceive this drop in quality and have a substantial deviation from human scores. We make similar observations across tasks that such metrics are not able to detect these simple perturbations.

Task-specific nuances are not captured. Our analysis also shows that existing metrics are not capable of addressing well known task specific grievances. For example, for the task of DG, it is known that many NLG systems generate generic responses which leads to poor engagement with the users. However, none of the metrics are sensitive to perturbations producing ‘generic responses’ such as *ok, thanks* or back-off responses such as *I’m sorry, can you repeat?*⁴ Similarly, for the task of QG, Nema and Khapra (2018) show that the *answerability* of a question is affected if we drop/replace question words or change named entities in the question. However, we find that most metrics (including the task specific QBLEU4/QROUGE) are not sensitive to such perturbations with very high deviation from human scores. On similar lines, ViLBERTScore which is a state of the art evaluation metric for IC is not sensitive to perturbations in gender, order of objects or attributes used for describing objects. This is of concern as many IC systems are known to produce generic captions containing genders, attributes and objects which are most prevalent in the training data. Similarly, for the task of D2T, where coverage and factual correctness are important we observe that most metrics are unable to detect perturbations which add extra/random text to the output or drop named entities (which often contain the most important information). Lastly, for AS it is important that an evaluation metric should penalise summaries which are not coherent or contain redundant sentences or do not have referential clarity. However, we observe that most metrics are not sensitive to perturbations which reorder the sentences or repeat sentences/phrases or replace nouns with pronouns (affecting referential clarity).

Different metrics have different skills. While no single metric is capable of detecting all types of perturbations, we observe that some metrics are

⁴A model that frequently produces generic responses is undesirable. Hence, the expert human annotators assigned a high penalty to this perturbation. Most of the task-specific metrics do use the context as an input whereas, the task-agnostic metrics do not.

more robust to certain perturbations. BLEURT and Moverscore are robust to jumbling of words, but BERTScore is not, revealing their differences in detecting *fluency*. Moverscore, BERTScore and the embedding-based metrics like Greedy Matching and Embedding Average are quite robust to simple transformations of converting numbers to corresponding words, which is an important criteria for the task of D2T, while BLEURT is relatively less robust to it. Similarly, while BERTScore performs poorly for many perturbations, it is able to respect alternative references, i.e., similar to humans, it does not drop its score when presented with alternative correct references from the dataset (last row in Figure 4a to 4f). An interesting observation from the IC task is that SoTA metrics like SPICE and ViBERTScore show a complementary behaviour on our set of perturbation criteria (third to last and last column in Figure 4f). This opens up interesting avenues for future research where different automatic metrics could be combined to take advantage of their relative strengths.

6 Related Work

Some of the related work, particularly the relevant datasets, human evaluation criteria, and automatic metrics were already discussed earlier and hence not covered again here. We refer the readers to two recent surveys (Sai et al., 2020b; Çelikyilmaz et al., 2020) for a detailed overview of automatic evaluation metrics as well as related work on criticising the use of automatic evaluation metrics. We mention a few such important works here. BLEU is one of the most widely analysed metric with several studies showing that it does not correlate well with human judgements for machine translation (Callison-Burch et al., 2006). This issue of poor correlations of metrics with human judgements has been reported on not just BLEU, but also on various other metrics, across several NLG tasks including Question Generation (Nema and Khapra, 2018), Data-to-Text generation (Dhingra et al., 2019), Dialogue generation (Liu et al., 2016), and Summarisation (Kryscinski et al., 2019). Apart from poor correlations, Kryscinski et al. (2019) criticize the automatic metrics for abstractive summarization since they don't check for factual inconsistencies in the summaries. Similarly Wiseman et al. (2017) discuss the lack of a reliable measurement of faithfulness in the context of Data-to-Text Generation. In case of dialogue, several n-gram-based

and embedding-based metrics have been shown to fall short in capturing the diversity of the valid responses (Liu et al., 2016; Sai et al., 2020a). The alternative of trained metrics, such as ADEM have been shown to be susceptible to adversarial attacks (Sai et al., 2019).

Similar to the main message of our work, some recent works have also called for a more robust evaluation of automatic evaluation metrics (Choshen and Abend, 2018; Mathur et al., 2020a). Ethayarajh and Jurafsky (2020) also critically examine the current approaches towards NLP leaderboards and point towards having multiple metrics along different dimensions such as fairness, efficiency, robustness, etc.

7 Conclusion

We conduct a large-scale study involving 6 tasks, 25 automatic evaluation metrics and 18 human evaluation criteria and observe that (i) different criteria such as fluency, coverage, etc are often not correlated and (ii) existing metrics have a low correlation with most criteria across different tasks. Based on these observations, we suggest an alternative framework for evaluating evaluation metrics which goes beyond computing correlations with the human scores for overall quality. More specifically, we propose perturbation templates which allow a more fine-grained evaluation of such metrics and help in understanding their strengths and more importantly their limitations. We hope that future work on designing evaluation metrics will use our perturbation checklist for evaluating the effectiveness of the proposed metric in assessing different relevant criteria.

Acknowledgements

We thank the annotators for helping us evaluate and annotate the perturbations. We thank the anonymous EMNLP-21 reviewers whose comments and feedback helped enhance the paper. We thank the Google India Ph.D. Fellowship Program and the Prime Minister's Fellowship Scheme for Doctoral Research for supporting Ananya Sai and Samsung IITM-Pravartak Undergraduate Fellowship for supporting Dev Yashpal Sheth. Finally, we thank the Robert Bosch Centre for Data Science and AI for supporting this work.

References

- Somak Aditya, Yezhou Yang, Chitta Baral, Cornelia Fermuller, and Yiannis Aloimonos. 2015. From images to sentences through scene description graphs using commonsense reasoning and knowledge. *arXiv preprint arXiv:1511.03292*.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. [SPICE: semantic propositional image caption evaluation](#). In *Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V*, volume 9909 of *Lecture Notes in Computer Science*, pages 382–398. Springer.
- Rafael E. Banchs, Luis F. D’Haro, and Haizhou Li. 2015. [Adequacy-fluency metrics: Evaluating MT in the continuous space model framework](#). *IEEE ACM Trans. Audio Speech Lang. Process.*, 23(3):472–482.
- 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.
- Chris Callison-Burch, Cameron Fordyce, Philipp Koehn, Christof Monz, and Josh Schroeder. 2007. [\(meta-\) evaluation of machine translation](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 136–158, Prague, Czech Republic. Association for Computational Linguistics.
- Chris Callison-Burch, Miles Osborne, and Philipp Koehn. 2006. [Re-evaluating the role of Bleu in machine translation research](#). In *11th Conference of the European Chapter of the Association for Computational Linguistics*, Trento, Italy. Association for Computational Linguistics.
- Thiago Castro Ferreira, Claire Gardent, Nikolai Ilinykh, Chris van der Lee, Simon Mille, Diego Moussallem, and Anastasia Shimorina. 2020. [The 2020 bilingual, bi-directional WebNLG+ shared task: Overview and evaluation results \(WebNLG+ 2020\)](#). In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 55–76, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Asli Çelikyilmaz, Elizabeth Clark, and Jianfeng Gao. 2020. [Evaluation of text generation: A survey](#). *CoRR*, abs/2006.14799.
- Leshem Choshen and Omri Abend. 2018. [Automatic metric validation for grammatical error correction](#). *CoRR*, abs/1804.11225.
- Julian Chow, Lucia Specia, and Pranava Madhyastha. 2019. [WMDO: Fluency-based word mover’s distance for machine translation evaluation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 494–500, Florence, Italy. 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 Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 2748–2760. Association for Computational Linguistics.
- Jan Deriu, Álvaro Rodrigo, Arantxa Otegi, Guillermo Echeгойen, Sophie Rosset, Eneko Agirre, and Mark Cieliebak. 2019. [Survey on evaluation methods for dialogue systems](#). *CoRR*, abs/1905.04071.
- Bhuvan Dhingra, Manaal Faruqui, Ankur P. Parikh, Ming-Wei Chang, Dipanjan Das, and William W. Cohen. 2019. [Handling divergent reference texts when evaluating table-to-text generation](#). In *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, pages 4884–4895. Association for Computational Linguistics.
- Kawin Ethayarajh and Dan Jurafsky. 2020. [Utility is in the eye of the user: A critique of NLP leaderboards](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 4846–4853. Association for Computational Linguistics.
- Alexander R. Fabbri, Wojciech Kryscinski, Bryan McCann, Caiming Xiong, Richard Socher, and Dragomir R. Radev. 2020. [Summeval: Re-evaluating summarization evaluation](#). *CoRR*, abs/2007.12626.
- Gabriel Forgues and Joelle Pineau. 2014. [Bootstrapping dialog systems with word embeddings](#). In *NeurIPS, modern machine learning and natural language processing workshop*.
- Markus Freitag, George F. Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *CoRR*, abs/2104.14478.
- Yang Gao, Wei Zhao, and Steffen Eger. 2020. [SUPER: towards new frontiers in unsupervised evaluation metrics for multi-document summarization](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 1347–1354. Association for Computational Linguistics.

- Sebastian Gehrmann, Tosin P. Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Aremu Anuoluwapo, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh D. Dhole, Wanyu Du, Esin Durmus, Ondrej Dusek, Chris Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, Harsh Jhamtani, Yangfeng Ji, Shailza Jolly, Dhruv Kumar, Faisal Ladhak, Aman Madaan, Mounica Maddela, Khyati Mahajan, Saad Mahamood, Bodhisattwa Prasad Majumder, Pedro Henrique Martins, Angelina McMillan-Major, Simon Mille, Emiel van Miltenburg, Moin Nadeem, Shashi Narayan, Vitaly Nikolaev, Rubungo Andre Niyongabo, Salomey Osei, Ankur P. Parikh, Laura Perez-Beltrachini, Niranjana Ramesh Rao, Vikas Raunak, Juan Diego Rodriguez, Sashank Santhanam, João Sedoc, Thibault Sellam, Samira Shaikh, Anastasia Shimorina, Marco Antonio Sobrevilla Cabezudo, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, and Jiawei Zhou. 2021. [The GEM benchmark: Natural language generation, its evaluation and metrics](#). *CoRR*, abs/2102.01672.
- Hany Hassan, Anthony Aue, Chang Chen, Vishal Chowdhary, Jonathan Clark, Christian Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, William Lewis, Mu Li, Shujie Liu, Tie-Yan Liu, Renqian Luo, Arul Menezes, Tao Qin, Frank Seide, Xu Tan, Fei Tian, Lijun Wu, Shuangzhi Wu, Yingce Xia, Dongdong Zhang, Zhirui Zhang, and Ming Zhou. 2018. [Achieving human parity on automatic chinese to english news translation](#). *CoRR*, abs/1803.05567.
- David M. Howcroft, Anya Belz, Miruna-Adriana Clinciu, Dimitra Gkatzia, Sadid A. Hasan, Saad Mahamood, Simon Mille, Emiel van Miltenburg, Sashank Santhanam, and Verena Rieser. 2020. [Twenty years of confusion in human evaluation: NLG needs evaluation sheets and standardised definitions](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 169–182, Dublin, Ireland. Association for Computational Linguistics.
- Ming Jiang, Qiuyuan Huang, Lei Zhang, Xin Wang, Pengchuan Zhang, Zhe Gan, Jana Diesner, and Jianfeng Gao. 2019. [TIGER: Text-to-image grounding for image caption evaluation](#). 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 2019, Hong Kong, China, November 3-7, 2019*, pages 2141–2152. Association for Computational Linguistics.
- Wojciech Kryscinski, Nitish Shirish Keskar, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. [Neural text summarization: A critical evaluation](#). 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 540–551, Hong Kong, China. Association for Computational Linguistics.
- Thomas K Landauer and Susan T Dumais. 1997. A solution to plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological review*, 104(2):211.
- Hwanhee Lee, Seunghyun Yoon, Franck Dernoncourt, Doo Soon Kim, Trung Bui, and Kyomin Jung. 2020. [ViBERTScore: Evaluating image caption using vision-and-language BERT](#). In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 34–39, 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.
- Chia-Wei Liu, Ryan Lowe, Iulian Serban, Michael 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, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2122–2132. The 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](#). *CoRR*, abs/1907.11692.
- Ryan Lowe, Michael Noseworthy, Iulian Vlad Serban, Nicolas Angelard-Gontier, Yoshua Bengio, and Joelle Pineau. 2017. [Towards an automatic turing test: Learning to evaluate dialogue responses](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pages 1116–1126. Association for Computational Linguistics.
- Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2020a. [Tangled up in BLEU: reevaluating the evaluation of automatic machine translation evaluation metrics](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4984–4997. Association for Computational Linguistics.
- Nitika Mathur, Johnny Wei, Markus Freitag, Qingsong Ma, and Ondrej Bojar. 2020b. [Results of the WMT20 metrics shared task](#). In *Proceedings of the Fifth Conference on Machine Translation, WMT@EMNLP 2020, Online, November 19-20, 2020*, pages 688–725. Association for Computational Linguistics.

- Preksha Nema and Mitesh M. Khapra. 2018. [Towards a better metric for evaluating question generation systems](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3950–3959. Association for Computational Linguistics.
- Jessica Ouyang, Serina Chang, and Kathy McKeown. 2017. [Crowd-sourced iterative annotation for narrative summarization corpora](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 2: Short Papers*, pages 46–51. 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.
- Maja Popovic. 2017. [chr++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation, WMT 2017, Copenhagen, Denmark, September 7-8, 2017*, pages 612–618. Association for Computational Linguistics.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [SQuAD: 100, 000+ questions for machine comprehension of text](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics.
- Ehud Reiter. 2018. [A structured review of the validity of BLEU](#). *Comput. Linguistics*, 44(3).
- Marco Túlio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with checklist](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 4902–4912. Association for Computational Linguistics.
- Vasile Rus and Mihai C. Lintean. 2012. [A comparison of greedy and optimal assessment of natural language student input using word-to-word similarity metrics](#). In *BEA@NAACL-HLT*.
- Ananya B. Sai, Mithun Das Gupta, Mitesh M. Khapra, and Mukundhan Srinivasan. 2019. [Re-evaluating ADEM: A deeper look at scoring dialogue responses](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 6220–6227. AAAI Press.
- Ananya B. Sai, Akash Kumar Mohankumar, Sidhartha Arora, and Mitesh M. Khapra. 2020a. [Improving dialog evaluation with a multi-reference adversarial dataset and large scale pretraining](#). *Trans. Assoc. Comput. Linguistics*, 8:810–827.
- Ananya B. Sai, Akash Kumar Mohankumar, and Mitesh M. Khapra. 2020b. [A survey of evaluation metrics used for NLG systems](#). *CoRR*, abs/2008.12009.
- Abigail See, Stephen Roller, Douwe Kiela, and Jason Weston. 2019. [What makes a good conversation? how controllable attributes affect human judgments](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 1702–1723. Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur P. Parikh. 2020. [BLEURT: learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 7881–7892. Association for Computational Linguistics.
- Shikhar Sharma, Layla El Asri, Hannes Schulz, and Jeremie Zumer. 2017. [Relevance of unsupervised metrics in task-oriented dialogue for evaluating natural language generation](#). *CoRR*, abs/1706.09799.
- Koustuv Sinha, Prasanna Parthasarathi, Jasmine Wang, Ryan Lowe, William L. Hamilton, and Joelle Pineau. 2020. [Learning an unreferenced metric for online dialogue evaluation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 2430–2441. Association for Computational Linguistics.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. [A study of translation edit rate with targeted human annotation](#). In *In Proceedings of Association for Machine Translation in the Americas*, pages 223–231.
- Chongyang Tao, Lili Mou, Dongyan Zhao, and Rui Yan. 2018. [RUBER: an unsupervised method for automatic evaluation of open-domain dialog systems](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 722–729. AAAI Press.
- Oleg V. Vasilyev, Vedant Dharnidharka, and John Bohannon. 2020. [Fill in the BLANC: human-free quality estimation of document summaries](#). *CoRR*, abs/2002.09836.

Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. *Cider: Consensus-based image description evaluation*. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015*, pages 4566–4575. IEEE Computer Society.

Sam Wiseman, Stuart M. Shieber, and Alexander M. Rush. 2017. *Challenges in data-to-document generation*. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pages 2253–2263. Association for Computational Linguistics.

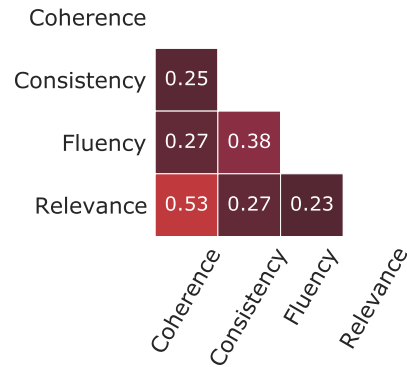
Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. *Bertscore: Evaluating text generation with BERT*. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

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 2019, Hong Kong, China, November 3-7, 2019*, pages 563–578. Association for Computational Linguistics.

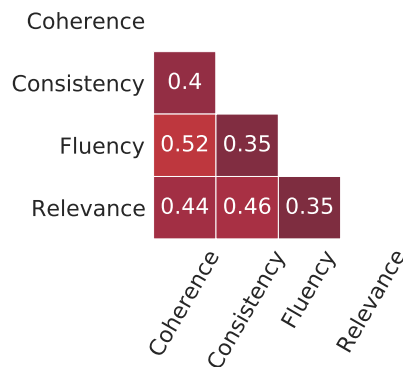
Wanzheng Zhu and Suma Bhat. 2020. *GRUEN for evaluating linguistic quality of generated text*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings, EMNLP 2020, Online Event, 16-20 November 2020*, pages 94–108. Association for Computational Linguistics.

A Criteria correlations

The pearson correlations among the criteria are presented in Figure 6. Most of the correlation ranges are similar for pearson correlation and kendall tau correlation, except for D2T task. We refer the studies on such correlations (Mathur et al., 2020a), discussing various points such as the influence of outliers and noisy points on the correlations. Additionally, we observe that the expertise of the annotators also influences the criteria-criteria correlations. In particular, we were able to study this in case of AS using the data released by Fabbri et al. (2020) containing both expert and crowdsourced annotations. From Figure 5, we observe that the scores by expert annotators have far lesser correlations amidst various criteria than the crowdsourced annotations.



(a) Expert annotations



(b) Crowdsourced annotations

Figure 5: Expert v/s Crowdsourced annotations

B Detailed examples to illustrate the implementation of the perturbation templates

Our perturbation templates mainly draw from the official github repository⁵ of the checklist paper and are also publicly available⁶. The implementation involves preprocessing with the help of tokenization, POS tagging, NER recognition, etc. Synonyms, antonyms, etc., are obtained with the help of WordNet framework. Additionally, the masked language model of RoBERTa is used to mask and predict replacements for the targeted words. For example, the application of the template for ‘dropping stop words’ involves the tokenization of the sentence using the NLTK word tokenizer as the first step. The list of tokens is compared with the set of stopwords provided by NLTK to filter out the stop words from the list of tokens. The modified sentence is then reconstructed using the string join function by iterating over the tokens in the modified list. Similarly, for the template of ‘changing the at-

⁵<https://github.com/marcotcr/checklist>

⁶<https://github.com/iitmnlp/EvalEval>

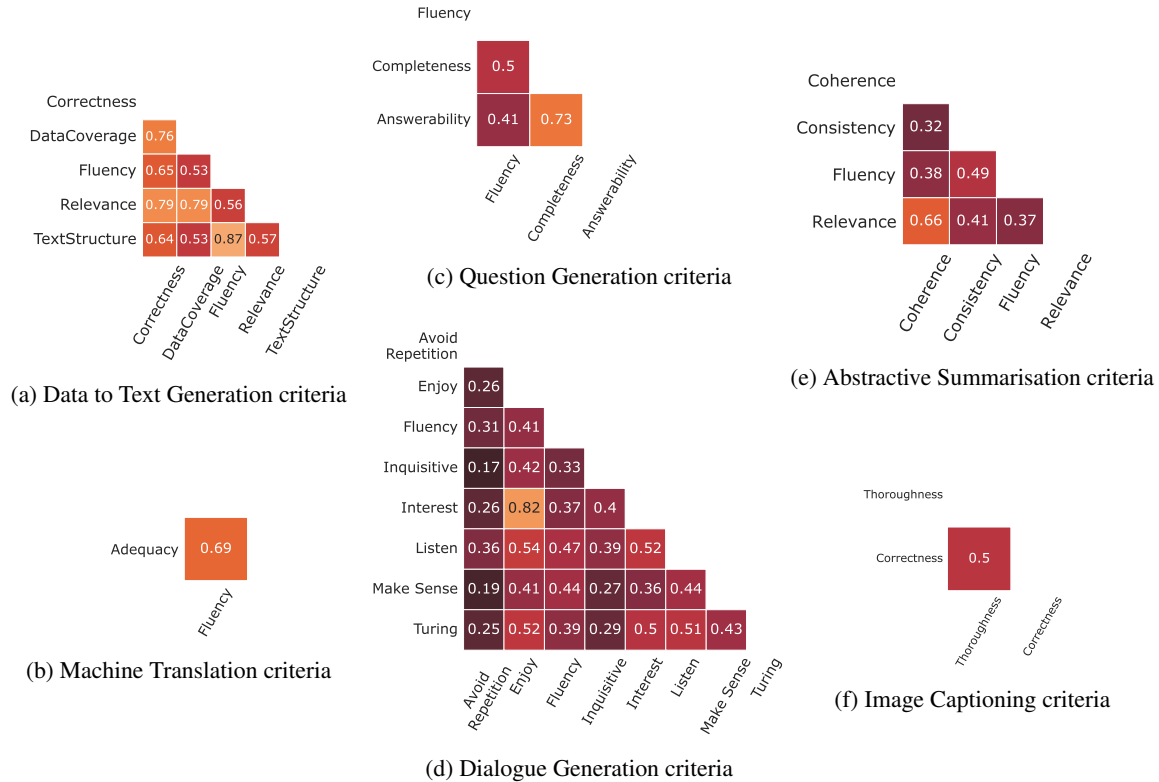


Figure 6: Correlations between criteria for 6 different tasks (MT, DG, AS, QG, D2T and IC)

tributes’ in case of image captioning, the sentence is first tokenized, then the adjectives are identified using part-of-speech tagging (again a functionality provided by NLTK). The list of ‘related words’ (i.e., hyponyms of hypernyms or ‘sibling words’) are obtained using WordNet framework. Unless the list returns empty, one of the entries in the list is used to replace the original adjective. In order to ‘change question to an assertive statement’, the question words (such as who, what, why, when, etc) are replaced with a ‘mask’ token and the ‘?’ character at the end is replaced with ‘.’ using string replace function. This modified sentence is then fed to RoBERTa model which generates different predictions to be used in place of the ‘mask’ token. One of the suggested words is used to form the modified assertive sentence. In case of perturbations involving dropping words, we additionally decide if we’re dropping stop words, adjectives, question words, etc in the particular perturbation and estimate the extent of effect it’ll have on each criteria. The perturbations of adding text, appends random words / phrases / sentences to a given text to account for not just the cases where there is missing information, but also cases where there is spurious wrong information, even if it accompanies

/ follows the correct version. The complete implementations / details of our perturbation templates are hosted publicly⁷.

Note that some of the perturbations cannot be applied to every sentence in the dataset. For example, the template of “changing names” cannot be applied if there are no named entities in a particular sentence. We hence shortlist only the successfully modified samples from the dataset for analysing the metrics’ performance on each perturbation.

C Perturbation Templates for various criteria

Table 5 contains the comprehensive list of perturbation templates used in our work.

D Correlations of Metrics with various Criteria (complete plots)

Figure 7 is a more comprehensive version of Figure 3. It shows the correlations of the complete set of metrics considered in this study with various criteria across different tasks.

⁷<https://github.com/iitmnlp/EvalEval>

Task	Criteria	Perturbation	Unmodified sentence	Perturbed sentence
All tasks	Fluency	Misplaced/missing punctuation	Could you let me know if I can meet him now or later ?	Could you let me know , if I can meet him now or later .
		Jumbling word order	We play badminton every evening.	We badminton every evening play.
		Subject-verb disagreement	He doesn't know how to bake.	He don't know how to bake.
		Dropping words (such as prepositions/articles, etc)	The bank is willing to approve the loan.	Bank willing to approve the loan.
		Spelling errors	Make the most of every opportunity presented to you.	Make the most of evry opportunity presented to you.
MT	Adequacy	Dropping out words or phrases	I was being followed.	I followed.
		Adding extra wrong information	This book is so inspiring.	This book is so inspiring, I forgot .
		Negation / antonyms	It will rain on Monday.	It will not rain on Monday.
		Repeat phrases	My relatives are in town.	My relatives are in town, my relatives .
AS	Informa-tiveness	Dropping words	Here is the no parking sign.	Here is the sign.
		Negation and antonyms	This book is so inspiring .	This book is so uninspiring .
	Flow / coherence	Use hyponyms to create misinformation	The girl my brother Andy met through MySpace turned out to be completely made up .	The girl my friend Andy met through MySpace turned out to be completely made up.
		Reorder sentences	The pandemic was spreading uncontrollably. Vaccines are being developed and tested rapidly.	Vaccines are being developed and tested rapidly. The pandemic was spreading uncontrollably.
		Non-Redundancy	Repeat sentences	My relatives are in town.
Referential clarity	Replace nouns by pronouns	The pandemic was spreading uncontrollably. Vaccines are being developed rapidly.	It was spreading uncontrollably. They are being developed rapidly.	
	QG	Change question word	When was he born?	What was he born?
Drop question word		When was he born ?	Was he born?	
Change question to assertive statement		Who is the director of Titanic?	The director of Titanic is James Cameron.	
Drop NLTk stop words		Who was not a world leader	Who world leader?	
Relevance	Mask a few words and predict	How could Tesla run the experiments ?	How could Tesla run to the beach ?	
	Perturb named entities, nouns, verbs	Why did Mary go to the school ?	Why did Raj go to the market ?	
Making sense	Negate a previous statement by same speaker	Bot: I enjoy having your daughter in my class. User: I'm glad to hear that.	I don't enjoy having your daughter in my class.	
	Add extra non-sensible text	Do you know where Dr. XYZ lives?	Yes, my father is my grandmother's father	
DG	Avoid repetition	Repeat previous utterances	Do you know Dr.XYZ?	Do you know where Dr. XYZ lives?
		Repeat phrases	I like ice creams	I like ice creams, ice creams
	Listening	Replace with "Can you repeat?"	I need to book a taxi	I'm sorry, can you repeat?
Relevance	Map random responses	I am new to coding .	I am scared of snakes .	
IC	Correctness	Change the order of objects	The man is standing in front of a tree	The tree is standing in front of man
		Change gender	Two girls are playing with a doll	Two boys are playing with a doll
	Thoroughness	Change attributes	A small boy playing with a red ball	A tall boy playing with a green ball
		Drop objects/noun	A small boy playing with a red ball	A small playing with a red
D2T	Correctness	Repeat (append) object	A lady riding a horse.	A lady riding a horse and a lady .
		Use hyponyms	Beethoven was a German musician	Beethoven was a German architect
		Change numbers	The cricketer was born in 1990 .	The cricketer was born in 1950 .
	Data Coverage	Negation / antonyms	The author of Harry Potter is J.K Rowling	The author of Harry Potter is not J.K Rowling.
		Drop phrases	A small boy playing with a red ball	A small boy playing with a
Relevance	Repeat phrases	Beethoven was a German musician	Beethoven was a German musician and German musician .	
	Random text	Beethoven was a German musician	The cricketer was born in 1990.	
	Perturb names	Phillips was a child prodigy.	James was a child prodigy.	
All tasks	Invariance	Replace with synonyms	The mangoes are delicious .	The mangoes are tasty .
		Contractions	We are going to embark on an adventure.	We're going to embark on an adventure.
		Expansions	There weren't any clear winners of the contest	There were not any clear winners of the contest.
		Numerals to words	Aron Ralston who was trapped for 127 hours.	Aron Ralston who was trapped for one hundred twenty seven hours.

Table 5: Perturbation templates targeting various criteria with examples. The blue highlights indicate the portions of the original sentence affected by the perturbation template. The red highlights indicate the changes in the modified sentence.

Fluency	0.36	0.35	0.33	0.31	0.26	0.34	0.35	0.4	0.29	0.24	0.26	0.34	0.19	0.34	0.38	0.4	0.35	0.33	0.36	0.25	0.25	0.24
Completeness	0.18	0.15	0.13	0.11	0.16	0.15	0.36	0.33	0.11	0.15	0.07	0.23	0.51	0.15	0.42	0.33	0.45	0.53	0.49	0.23	0.21	0.23
Answerability	0.18	0.15	0.13	0.1	0.08	0.19	0.25	0.28	0.12	0.19	0.16	0.16	0.34	0.13	0.3	0.25	0.31	0.37	0.37	0.21	0.18	0.21
	Bleu_1	Bleu_2	Bleu_3	Bleu_4	ROUGE-4	ROUGE-L	METEOR	chrF++	Eavg	VE	GM	WMD	SMS	GRUEN	BertScore-f	BertScore-p	BertScore-r	Movescore	BLEURT	QBLEU1	QBLEU4	QROUGE-L

(a) Question Generation (QG)

Coherence	0.03	0.02	0.02	0.02	0.11	0.13	0.11	0.12	0.04	0.09	0.01	0.11	0.12	0.26	0.13	0.09	0.13	0.12	0.1	0.14	0.11	0.17	0.08
Consistency	0.07	0.04	0.04	0.03	0.13	0.11	0.12	0.12	0.01	0.01	0.05	0.12	0.16	0.04	0.12	0.08	0.12	0.11	0.07	0.12	0.1	0.24	0.15
Fluency	0.04	0.03	0.03	0.03	0.07	0.04	0.07	0.06	0.03	0.03	0.01	0.1	0.1	0.13	0.12	0.12	0.08	0.09	0.12	0.05	0.03	0.14	0.06
Relevance	0.04	0.03	0.02	0.02	0.17	0.22	0.21	0.24	0.01	0.08	0.01	0.13	0.2	0.17	0.2	0.12	0.21	0.21	0.2	0.25	0.2	0.22	0.15
	Bleu_1	Bleu_2	Bleu_3	Bleu_4	ROUGE-4	ROUGE-L	METEOR	chrF++	Eavg	VE	GM	WMD	SMS	GRUEN	BertScore-f	BertScore-p	BertScore-r	Movescore	BLEURT	Rouge-su * r	Rouge-w r	SUPERT	BLANC

(b) Abstractive Summarization (AS)

Correctness	0.2	0.2	0.18	0.16	0.19	0.23	0.24	0.24	0.2	0.1	0.17	0.16	0.21	0.21	0.18	0.26	0.24	0.25	0.32	0.28	0.07
DataCoverage	0.18	0.17	0.15	0.13	0.16	0.18	0.23	0.24	0.17	0.07	0.14	0.12	0.19	0.19	0.19	0.23	0.18	0.24	0.28	0.26	0.05
Fluency	0.2	0.2	0.19	0.17	0.2	0.26	0.22	0.23	0.21	0.09	0.17	0.16	0.2	0.2	0.17	0.27	0.26	0.25	0.35	0.28	0.05
Relevance	0.16	0.16	0.15	0.13	0.15	0.2	0.19	0.2	0.17	0.06	0.13	0.13	0.17	0.18	0.15	0.21	0.2	0.2	0.26	0.23	0.05
TextStructure	0.18	0.19	0.17	0.15	0.19	0.25	0.21	0.21	0.19	0.07	0.15	0.14	0.18	0.18	0.15	0.25	0.24	0.23	0.33	0.25	0.04
	Bleu-1	Bleu-2	Bleu-3	Bleu-4	ROUGE-L	TER	METEOR	chrF++	Sacrebleu	EA	VE	GM	WMD	WYDO	SMS	BERTscore F1	BERTscore P	BERTscore R	BLEURT	Movescore	PARENT

(c) Data to Text Generation (D2T)

Correctness	0.31	0.3	0.27	0.26	0.27	0.35	0.28	0.27	0.37	0.33	0.24	0.34	0.37	0.33	0.36	0.39	0.25	0.38	0.37	0.38	0.37	0.38
Thoroughness	0.29	0.28	0.25	0.24	0.25	0.31	0.24	0.25	0.33	0.29	0.2	0.3	0.33	0.3	0.33	0.35	0.23	0.33	0.32	0.33	0.33	0.34
	Bleu_1	Bleu_2	Bleu_3	Bleu_4	ROUGE-L	METEOR	chrF++	Eavg	VE	GM	WMD	SMS	BertScore-p	BertScore-r	BertScore-f	BLEURT	TER	SPACE	TICEF	VILBERTScore**p	VILBERTScore**r	VILBERTScore**f

(d) Image Captioning

MaUde	0.25	0.14	0.09	0.06	0.19	0.2	0.25	0.26	0.16	0.2
BERT NLI	0.29	0.1	0.11	0.16	0.15	0.23	0.24	0.18	0.07	0.19
Inferent	0.18	0.04	0.05	0.09	0.1	0.15	0.18	0.22	0.14	0.13
DEB	0.18	0.03	0.02	0.1	0.04	0.07	0.01	0.1	0.13	0.1
RUBER	0.22	0.12	0.09	0.11	0.11	0.19	0.21	0.2	0.16	0.17
	Enjoy	Avoid repetition	Turning	Inquisitiveness	Makes sense	Interest	Persona guess	Fluency	Listen	Average

(e) Dialogue Generation (DG)

Figure 7: Correlations of metrics with different criteria