

# Curious Case of Language Generation Evaluation Metrics: A Cautionary Tale

Ozan Caglayan<sup>1</sup>, Pranava Madhyastha<sup>1</sup> and Lucia Specia<sup>1,2,3</sup>

Imperial College London<sup>1</sup>, University of Sheffield<sup>2</sup>, ADAPT - Dublin City University<sup>3</sup>

o.caglayan@ic.ac.uk, pranava@ic.ac.uk, l.specia@ic.ac.uk

## Abstract

Automatic evaluation of language generation systems is a well-studied problem in Natural Language Processing. While novel metrics are proposed every year, a few popular metrics remain as the *de facto* metrics to evaluate tasks such as image captioning and machine translation, despite their known limitations. This is partly due to ease of use, and partly because researchers expect to see them and know how to interpret them. In this paper, we urge the community for more careful consideration of how they automatically evaluate their models by demonstrating important failure cases on multiple datasets, language pairs and tasks. Our experiments show that metrics (i) usually prefer system outputs to human-authored texts, (ii) can be insensitive to correct translations of rare words, (iii) can yield surprisingly high scores when given a single sentence as system output for the entire test set.

## 1 Introduction

Human assessment is the best practice at hand for evaluating language generation tasks such as machine translation (MT), dialogue systems, visual captioning and abstractive summarisation. In practice, however, we rely on automatic metrics which compare system outputs to human-authored references. Initially proposed for MT evaluation, metrics such as BLEU (Papineni et al., 2002) and METEOR (Denkowski and Lavie, 2014) are increasingly used for other tasks, along with more task-oriented metrics such as ROUGE (Lin, 2004) for summarisation and CIDER (Vedantam et al., 2015) for visual captioning.

Reiter and Belz (2009) remark that it is *not sufficient* to conclude on the usefulness of a natural language generation system’s output by solely relying on metrics that quantify the similarity of the output to human-authored texts. Previous criticisms concerning automatic metrics corroborate this perspective to some degree. To cite a few, in the context of MT, Callison-Burch et al. (2006) state that human judgments may not correlate with BLEU and an increase in BLEU does not always indicate an improvement in quality. Further, Mathur et al. (2020) challenge the stability of the common practices measuring correlation between metrics and human judgments, the standard approach in the MT community, and show that they may be severely impacted by outlier systems and the sample size. In the context of image captioning, Wang and Chan (2019) claim that the consensus-based evaluation protocol of CIDER actually penalises output diversity. Similar problems have also been discussed in the area of automatic summarisation with respect to ROUGE (Schluter, 2017). Nevertheless, automatic metrics like these are a necessity and remain popular, especially given the increasing number of open evaluation challenges.

In this paper, we further probe BLEU, METEOR, CIDER-D and ROUGE<sub>L</sub> metrics (§ 2) that are commonly used to quantify progress in language generation tasks. We also probe a recently proposed contextualised embeddings-based metric called BERTSCORE (Zhang et al., 2020). We first conduct *leave-one-out* average scoring with multiple references and show that, counter-intuitively, metrics tend to reward system outputs more than human-authored references (§ 3.1). A system output perturbation experiment further highlights how metrics penalise errors in extremely frequent *n*-grams (§ 3.2.1) while they are quite insensitive to errors in rare words (§ 3.2.2). The latter makes it hard to understand whether a model

---

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: <http://creativecommons.org/licenses/by/4.0/>.

is better than another in its ability to handle rare linguistic phenomena correctly. Finally, we design (§ 3.3) an adversary that seeks to demonstrate metrics’ preference for frequent  $n$ -grams by using a single training set sentence as system output for the entire test set. We observe strikingly high scores, such as the sentence “*A man is playing*” obtaining a BLEU score of 30.6, compared to 47.9 of a strong model, on a captioning task. We hope that our observations in this paper will lead the community towards formulation of better metrics and evaluation protocols in the future (§ 4).

## 2 Automatic Metrics

In this section, we briefly describe the metrics we use in our experiments. To compute BLEU, METEOR, CIDER and ROUGE-L, we provide pre-tokenised hypotheses and references to `coco-caption` utility<sup>1</sup>. For BERTSCORE, we use its official release<sup>2</sup>.

- Initially proposed for MT evaluation, **BLEU** is a prevalent metric based on  $n$ -gram matches between the candidate and reference sentences. The final score is the geometric mean of  $n$ -gram *precisions* with a brevity penalty to penalise outputs shorter than references.
- In contrast, **METEOR** is a *recall*-based metric which uses explicit word alignments between candidate and references, allowing for exact matching, synonymy, and stemming. A fragmentation penalty rewards longer and fewer chunks of contiguously aligned tokens. The final score is computed between the *best* scoring reference and the candidate sentence.
- Another *recall*-based metric is **ROUGE-L** which measures longest common sub-sequences between the candidate and references, *i.e.* a set of shared words with similar order even if not contiguous.
- **CIDER** is a consensus-based captioning metric which uses term frequency inverse document frequency (TF-IDF) for  $n$ -gram weighting over all references. The final score is the average cosine similarity between the candidate sentence and references. We use the popular variant CIDER-D, which integrates a length-based Gaussian penalty and clipping.
- Finally, **BERTSCORE** computes a token level similarity for each candidate token against each token in the reference sentence, using contextual BERT embeddings (Zhang et al., 2020). We report the F-score, *i.e.* the harmonic mean of the precision and the recall of maximal per-token cosine similarities between the candidate and the reference(s).

## 3 Experiments

This section describes our probing experiments and results<sup>3</sup> for different tasks and language pairs.

### 3.1 Machine vs. human-authored texts

We first explore the extent to which metrics reward human-authored texts. We hypothesise that metrics should prefer human-authored texts to machine-produced texts, as the former is considered to be the ‘ground truth’ reference. For a given **multi-reference** metric  $M$ , REF-VS-REF *leave-one-out*<sup>4</sup> average  $L_M$  is computed across  $C$  human-authored reference corpora  $\{R_1, \dots, R_C\}$ . At each iteration  $i$ , the reference  $R_i$  is taken as the candidate and is evaluated against the held-out  $C-1$  references:

$$L_M = \frac{1}{C} \sum_{i=1}^C M(\text{SYS}=R_i, \text{REFS}=\{R_j : j \neq i\}) \quad (1)$$

Same principle holds when evaluating a trained system, except that SYS would now represent the system outputs. The metric computations inside the sum would still be using  $C-1$  references at each iteration.

<sup>1</sup><https://github.com/tylin/coco-caption>

<sup>2</sup>Version 0.3.4 `hug_trans 2.11.0`, with default settings except `rescale-with-baseline` is enabled.

<sup>3</sup>BLEU, METEOR, CIDER, ROUGE, BERTSCORE may sometimes be abbreviated as BL, MT, CR, RG, BS, for readability.

<sup>4</sup>Our approach differs from Vinyals et al. (2017) in that we go beyond BLEU1.

	BL	MT	CR	RG	Bs
NIST-2008 Chinese→English ( $C = 4$ )					
RvR	38.2 ± 2.2	<b>33.1</b> ± .9	<b>2.20</b> ± .10	59.5 ± 1.8	<b>0.69</b> ± .020
SvR	<b>41.8</b> ± .7	32.6 ± .3	2.13 ± .07	59.5 ± .5	0.68 ± .003
COCO Image→English ( $C = 5$ )					
RvR	19.5 ± .3	24.2 ± .1	0.88 ± .01	46.7 ± .3	0.55 ± .003
SvR	<b>30.3</b> ± .8	<b>25.1</b> ± .3	<b>1.05</b> ± .02	<b>53.0</b> ± .5	<b>0.61</b> ± .004
VATEX Video→English ( $C = 10$ )					
RvR	24.1 ± .2	<b>24.9</b> ± .1	<b>0.63</b> ± .01	48.8 ± .1	0.57 ± .002
SvR	<b>30.6</b> ± .1	23.0 ± .1	0.59 ± .01	<b>51.4</b> ± .1	<b>0.60</b> ± .000
VATEX Video→Chinese ( $C = 10$ )					
RvR	17.3 ± 1.5	22.1 ± .7	0.42 ± .01	45.8 ± 1.4	0.56 ± .015
SvR	<b>25.3</b> ± .7	<b>24.5</b> ± .3	<b>0.44</b> ± .01	<b>50.8</b> ± .5	<b>0.59</b> ± .005

Table 1: Leave-one-out averages for REF vs. REF (RvR) and SYS vs. REF (SvR) evaluations.

→UNK	BL	MT	CR	RG	Bs
–	19.95	24.41	0.896	47.4	0.551
people (0.61%)	19.53 ↓ .42	24.24 ↓ .17	0.878 ↓ .018	47.0 ↓ .4	0.547 ↓ .004
standing (0.88%)	19.45 ↓ .50	24.12 ↓ .29	0.869 ↓ .027	47.0 ↓ .4	0.542 ↓ .009
sitting (1.06%)	19.11 ↓ .84	24.02 ↓ .39	0.859 ↓ .037	47.0 ↓ .4	0.540 ↓ .011
man (1.21%)	19.11 ↓ .84	23.66 ↓ .75	0.859 ↓ .037	46.6 ↓ .8	0.539 ↓ .012
a (16.7%)	9.20 ↓ 10.75	20.03 ↓ 4.38	0.490 ↓ .406	34.8 ↓ 12.6	0.480 ↓ .071
RANDOM (16.7%)	13.63 ↓ 6.32	19.58 ↓ 4.83	0.500 ↓ .396	41.7 ↓ 5.7	0.413 ↓ .138

Table 2: Frequent 1-gram substitutions: % is the percentage of tokens mapped to UNK.

**Tasks.** We explore two visual captioning tasks, namely, COCO English image captioning (Chen et al., 2015) and the bilingual VATEX video captioning dataset (Wang et al., 2019). For the former, we use a state-of-the-art captioning model (Lu et al., 2018), and evaluate the 5k sentences in the test split (Karpathy and Fei-Fei, 2015). For English and Chinese VATEX, we train neural baselines and evaluate on the *dev* set which contains 3k videos. Finally, we also experiment with the Chinese→English NIST 2008 (MT08) test set which contains 1357 sentences, and translate it using *Google Translate*.

**Results.** Table 1 shows that human-authored references obtain relatively worse scores than systems in general. BLEU exhibits the biggest differences, with almost 11 points for COCO and 3.5 points for the NIST MT experiment. The results are consistent across different datasets and languages except for English VATEX and NIST-2008, which score slightly better than system outputs in METEOR and CIDER. Overall, the results suggest that the metrics do not necessarily reflect the quality of the generated captions or translations, as scores for human-authored texts are far from the upper bounds of the metrics (*e.g.* 100 for BLEU). Therefore, we suggest that this type of evaluation should not be used to draw conclusions on ‘human parity’, as previously done in some studies (Xu et al., 2015; Vinyals et al., 2017).

### 3.2 $N$ -gram perturbations

In this section, we propose two perturbation experiments to investigate the sensitivity of metrics to frequent and infrequent phenomena. In both experiments, the first reference corpus  $R_1$  of COCO test set is considered to be the output of a hypothetical system from which perturbed versions are created. Each version is then evaluated against the four remaining references  $\{R_2, \dots, R_5\}$  of the COCO test set. Although we target unigrams for perturbation, they implicitly effect higher order  $n$ -grams as well.

#### 3.2.1 Frequent unigram perturbation

We create five independent variants of  $R_1$  by substituting the most common words ‘*people*’, ‘*standing*’, ‘*sitting*’, ‘*man*’ and ‘*a*’ with the unknown token UNK. Table 2 shows that metrics react quite conservatively to this attack, until very high substitution percentages are reached. For example, substituting ‘*people*’ yields a drop of less than 0.4 points for BLEU, METEOR and ROUGE, a potentially uninteresting drop which can easily be overlooked during model development.

Although repeatedly missing indefinite articles ‘*a*’ is *semantically* much less critical than missing ‘*people*’, all metrics aggressively penalise the former, simply because the function word ‘*a*’ is extremely frequent: with 53.8% relative drop, BLEU is the most affected metric, whereas METEOR and BERTSCORE are more robust with relative drops of 17.9% and 12.9%, respectively. To find out more, we randomly substitute content words from  $R_1$  (RANDOM in Table 2), until we reach the same percentage (16.7%) as

$\geq T$	UNK	BL	MT	CR	RG	BS
30	1.44%	19.80	24.25	0.882	47.2	0.543
20	1.07%	19.84	24.30	0.887	47.3	0.545
10	0.67%	19.89	24.34	0.891	47.3	0.548
5	0.41%	19.91	24.37	0.893	47.3	0.549
1	0%	19.96	24.41	0.896	47.4	0.551
woman → man		19.63	24.15	0.889	47.0	0.549

Table 3: Insensitivity to infrequent words:  $T$  is the occurrence threshold used to obtain the vocabularies.

	TRAIN	TEST	REFS
MSVD	39K	670	41
DAILYDIALOG	87K	6.7K	5
FLICKR30K	149K	1K	5
VATEX	259K	3K	10
COCO	550K	5K	5
MT	1.6M	1.3K	4

Table 4: Dataset statistics: unique counts are given for training set sentences.

‘a’. The results are interesting: BLEU (9.20→13.63) and ROUGE (34.8→41.7) show substantial increases when compared to dropping ‘a’, supporting our *semantic* conjecture above. However, the decreases in METEOR (20.03→19.58) and BERTSCORE (0.480→0.413) highlight the sensitivity of these metrics to semantics rather than pure frequencies. The continuous and contextual nature of BERTSCORE seems to be an advantage here.

### 3.2.2 Insensitivity to infrequent constructions

We now approach to the problem from the other end to explore how sensitive metrics are to a set of rarely occurring words. We conjecture that this will provide insight into how much the metrics reward a hypothetical model that systematically translates a set of rare words correctly. Specifically, we create four variants of  $R_1$ , with each one substituting a particular set of words by UNK, based on training set frequencies.<sup>5</sup> Table 3 shows that even the most aggressively short-listed hypothetical model ( $T = 30$ ) obtains marginally worse scores than the full vocabulary model ( $T = 1$ ). As a concrete example, the last row shows the non-substantial impact of systematically replacing each occurrence of ‘woman’ (0.54% of all unigrams) with ‘man’ (1.21% of all unigrams).

**Discussion.** Although it is hard to interpret the magnitude of differences observed in both experiments, our empirical findings highlight interesting cases. Overall, the conclusions of our perturbation experiments are in line with the concurrent work of Mathur et al. (2020) which suggests that important conclusions, such as comparative judgments about systems, should not be drawn based only on small changes in automatic metrics.

### 3.3 Single representative sentence

Following the observations from the previous experiments, we search over the training set for a *single representative sentence* (SS) which maximises test set BLEU<sup>6</sup> when used as a system output for *every* test set instance.<sup>7</sup> We explore **tasks** and **datasets** which include the ones previously introduced (§ 3.1). For visual captioning, we add MSVD (Chen and Dolan, 2011), a widely known dataset of 1,970 videos, with up to 41 English captions per video. For image captioning, we use English Flickr30k (Young et al., 2014), and the STAIR (Yoshikawa et al., 2017) dataset which provides Japanese captions for COCO images. We also explore the multi-turn dialogue dataset DailyDialog (Li et al., 2017) which contains conversations that cover 10 different daily life topics, and its multi-reference test set (Gupta et al., 2019). Table 4 summarises the statistics about the datasets explored for this experiment.

Table 5 draws a comparison between the scores obtained for the retrieved single sentences, baselines and state-of-the-art systems when available. We observe that: (i) MSVD exhibits the highest scores with 30.6 BLEU and 23.4 METEOR, the latter being very close to a strong captioning baseline (Venugopalan et al., 2015), (ii) the SS scores for the DAILYDIALOG are surprisingly on par with a recent baseline and a state-of-the-art system, and (iii) CIDER is more robust against the single sentence adversary as more references seem to affect its internal consensus-based scoring. Although BERTSCORE does not explicitly

<sup>5</sup>This simulates the construction of a training vocabulary by only retaining tokens that occur at least  $T$  times *i.e.* short-listing.

<sup>6</sup>Maximising other metrics could lead to finding other sentences and therefore different results.

<sup>7</sup>We note that the retrieved sentences almost never occur in test set references.

DATASET	METHOD	BL1	BL4	MT	CR	RG	BS
COCO-EN	<i>“A pizza is sitting on top of a table”</i>	46.4	9.1	10.5	0.07	34.8	0.373
	Neural Baby Talk (Lu et al., 2018)	75.2	34.4	26.5	1.06	55.5	0.628
FLICKR-EN	<i>“A man in a blue shirt standing in front of a building”</i>	46.2	12.2	11.6	0.07	32.7	–
	Hard Attention (Xu et al., 2015)	66.9	19.9	18.5	–	–	–
	Neural Baby Talk (Lu et al., 2018)	69.0	27.1	21.7	0.58	–	–
MSVD-EN	<i>“A man is playing”</i>	73.8	30.6	23.4	0.15	59.7	–
	(Venugopalan et al., 2015, FLOW)	–	–	24.3	–	–	–
	State-of-the-art (Aafaq et al., 2019)	–	47.9	35.0	0.78	71.5	–
VATEX-EN	[test] <i>“A man is working out in front of a group of people”</i>	54.2	14.0	13.6	0.05	34.6	–
	[test] Baseline (Wang et al., 2019)	71.3	28.5	21.6	0.45	47.0	–
	[test] Leaderboard Winner	81.9	39.1	25.8	0.73	53.3	–
	[dev] <i>“A man is working out in front of a group of people”</i>	57.3	12.9	14.5	0.05	38.9	0.500
	[dev] Baseline (Wang et al., 2019)	76.3	32.0	23.3	0.59	52.1	0.605
DDIALOG-EN	<i>“Where do you want to go?”</i>	36.5	3.8	10.9	0.04	24.3	0.225
	Seq2Seq (Gupta et al., 2019)	–	4.5	9.7	–	29.3	–
	SoTA (Gupta et al., 2019, HRED)	45.8	6.2	11.1	0.10	32.9	0.289
COCO-JA	<i>“A person who is skiing stands”</i>	51.5	16.3	17.8	0.08	37.2	–
	EN caption + MT (Yoshikawa et al., 2012)	56.5	12.7	–	0.32	44.9	–
	End-to-end (Yoshikawa et al., 2012)	76.3	38.5	–	0.83	55.3	–
VATEX-ZH	[test] <i>“In a room, a man in black is dancing”</i>	70.8	21.7	27.7	0.06	48.6	–
	[test] Baseline (Wang et al., 2019)	74.5	24.8	29.4	0.35	51.6	–
	[test] Leaderboard Winner	83.0	32.6	32.5	0.64	56.7	–
	[dev] <i>“In a room, a man in black is dancing”</i>	68.2	19.2	22.5	0.07	45.8	0.531
	[dev] Baseline (Wang et al., 2019)	75.4	26.8	25.0	0.44	51.7	0.597

Table 5: Comparison of single sentence scores to several baselines and state-of-the-art: “*Single representative sentences are in italics.*” English translations are provided for non-English tasks.

rely on  $n$ -gram statistics, it also exhibits surprisingly high scores for the SS setting. This contradicts the salient claims regarding the utility of the metric for model selection (Zhang et al., 2020).

## 4 Discussion

In this work, we explore cases where commonly used language generation evaluation metrics exhibit *counter-intuitive* behaviour. Although the main goal in language generation tasks is to generate ‘human-quality’ texts, our analysis in § 3.1 shows that metrics have a preference towards machine-generated texts rather than human references. Our perturbation experiments in § 3.2.2 highlight potential insensitivity of metrics to lexical changes in infrequent  $n$ -grams. This is a major concern for tasks such as multimodal machine translation (Specia et al., 2016) or pronoun resolution in MT (Guillou et al., 2016), where the metrics are expected to capture lexical changes which are due to rarely occurring linguistic ambiguities. We believe that targeted probes (Isabelle et al., 2017) are much more reliable than sentence or corpus level metrics for such tasks. Finally, we reveal that metrics tend to produce unexpectedly high scores when each test set hypothesis is set to a particular training set sentence, which can be thought of as finding the *sweet spot* of the corpus (§ 3.3). Therefore, we note that a high correlation between metrics and human judgments is not sufficient to characterise the reliability of a metric. The latter probably requires a thorough exploration and mitigation of adversarial cases such as the proposed *single sentence* baseline.

## Acknowledgments

We thank the anonymous ACL-IJCNLP and COLING reviewers whose suggestions helped refine the paper. The authors received funding from the MultiMT (H2020 ERC Starting Grant No. 678017) and MMVC (Newton Fund Institutional Links Grant, ID 352343575) projects.

## References

- Nayyer Aafaq, Naveed Akhtar, Wei Liu, Syed Zulqarnain Gilani, and Ajmal Mian. 2019. Spatio-temporal dynamics and semantic attribute enriched visual encoding for video captioning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12487–12496.
- Chris Callison-Burch, Miles Osborne, and Philipp Koehn. 2006. Re-evaluation the role of Bleu in machine translation research. In *11th Conference of the European Chapter of the Association for Computational Linguistics*, Trento, Italy, April. Association for Computational Linguistics.
- David Chen and William Dolan. 2011. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 190–200, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, and C. Lawrence Zitnick. 2015. Microsoft COCO Captions: Data Collection and Evaluation Server. *arXiv preprint arXiv:1504.00325*, April.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 376–380, Baltimore, Maryland, USA, June. Association for Computational Linguistics.
- Liane Guillou, Christian Hardmeier, Preslav Nakov, Sara Stymne, Jörg Tiedemann, Yannick Versley, Mauro Cettolo, Bonnie Webber, and Andrei Popescu-Belis. 2016. Findings of the 2016 WMT shared task on cross-lingual pronoun prediction. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 525–542, Berlin, Germany, August. Association for Computational Linguistics.
- Prakhar Gupta, Shikib Mehri, Tiancheng Zhao, Amy Pavel, Maxine Eskenazi, and Jeffrey Bigham. 2019. Investigating evaluation of open-domain dialogue systems with human generated multiple references. In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 379–391, Stockholm, Sweden, September. Association for Computational Linguistics.
- Pierre Isabelle, Colin Cherry, and George Foster. 2017. A challenge set approach to evaluating machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2486–2496, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 3128–3137.
- Yanran Li, Hui Su, Xiaoyu Shen, Wenjie Li, Ziqiang Cao, and Shuzi Niu. 2017. DailyDialog: A manually labelled multi-turn dialogue dataset. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 986–995, Taipei, Taiwan, November. Asian Federation of Natural Language Processing.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain, July. Association for Computational Linguistics.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2018. Neural baby talk. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7219–7228.
- Nitika Mathur, Timothy Baldwin, and Trevor Cohn. 2020. 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*, pages 4984–4997, Online, July. 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, July. Association for Computational Linguistics.
- Ehud Reiter and Anja Belz. 2009. An investigation into the validity of some metrics for automatically evaluating natural language generation systems. *Computational Linguistics*, 35(4):529–558.
- Natalie Schluter. 2017. The limits of automatic summarisation according to ROUGE. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 41–45, Valencia, Spain, April. Association for Computational Linguistics.

- Lucia Specia, Stella Frank, Khalil Sima'an, and Desmond Elliott. 2016. A shared task on multimodal machine translation and crosslingual image description. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 543–553, Berlin, Germany, August. Association for Computational Linguistics.
- 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.
- Subhashini Venugopalan, Marcus Rohrbach, Jeffrey Donahue, Raymond Mooney, Trevor Darrell, and Kate Saenko. 2015. Sequence to sequence-video to text. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4534–4542.
- O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. 2017. Show and Tell: Lessons Learned from the 2015 MSCOCO Image Captioning Challenge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(4):652–663, April.
- Qingzhong Wang and Antoni B. Chan. 2019. Describing Like Humans: On Diversity in Image Captioning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, June.
- X. Wang, J. Wu, J. Chen, L. Li, Y. Wang, and W. Y. Wang. 2019. VaTeX: A Large-Scale, High-Quality Multilingual Dataset for Video-and-Language Research. In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4580–4590.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2048–2057. JMLR Workshop and Conference Proceedings.
- Katsumasa Yoshikawa, Ryu Iida, Tsutomu Hirao, and Manabu Okumura. 2012. Sentence compression with semantic role constraints. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 349–353, Jeju Island, Korea, July. Association for Computational Linguistics.
- Yuya Yoshikawa, Yutaro Shigeto, and Akikazu Takeuchi. 2017. STAIR captions: Constructing a large-scale Japanese image caption dataset. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 417–421, Vancouver, Canada, July. Association for Computational Linguistics.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.
- 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*.