

Re-evaluating Automatic Metrics for Image Captioning

Mert Kilickaya, Aykut Erdem, Nazli Ikizler-Cinbis, and Erkut Erdem

Hacettepe University Computer Vision Lab

Dept. of Computer Engineering, Hacettepe University, Ankara, TURKEY

kilickayamert@gmail.com, {aykut, nazli, erkut}@cs.hacettepe.edu.tr

Abstract

The task of generating natural language descriptions from images has received a lot of attention in recent years. Consequently, it is becoming increasingly important to evaluate such image captioning approaches in an automatic manner. In this paper, we provide an in-depth evaluation of the existing image captioning metrics through a series of carefully designed experiments. Moreover, we explore the utilization of the recently proposed Word Mover's Distance (WMD) document metric for the purpose of image captioning. Our findings outline the differences and/or similarities between metrics and their relative robustness by means of extensive correlation, accuracy and distraction based evaluations. Our results also demonstrate that WMD provides strong advantages over other metrics.

1 Introduction

There has been a growing interest in research on integrating vision and language in natural language processing and computer vision communities. As one of the key problems in this emerging area, image captioning aims at generating natural descriptions of a given image (Bernardi et al., 2016). This is a challenging problem since it requires the ability to not only understand the visual content, but also to generate a linguistic description of that content. In this regard, it can be framed as a machine translation task where the source language denotes the visual domain and the target language is a specific language such as English. The recently proposed deep image captioning studies follow this interpretation and model the process

via an encoder-decoder architecture (Vinyals et al., 2015; Xu et al., 2015; Karpathy and Fei-Fei, 2015; Jia et al., 2015). These approaches have attained considerable success in the recent benchmarks such as FLICKR8K (Hodosh et al., 2013), FLICKR30K (Young et al., 2014) and MS COCO (Lin et al., 2014) as compared to the earlier techniques which explicitly detect objects and generate descriptions by using surface realization techniques (Kulkarni et al., 2013; Li et al., 2011; Elliott and Keller, 2013).

With the size of the benchmark datasets becoming larger and larger, evaluating image captioning models has become increasingly important. Human-based evaluations become obsolete as they are costly to acquire and, more importantly, not repeatable. Automatic evaluation metrics are employed as an alternative to human evaluation in both developing new models and comparing them against the state-of-the-art. These metrics compute a score that indicates the similarity/dissimilarity between an automatically generated caption and a number of human-written reference (gold standard) descriptions.

Some of these automatic metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Banerjee and Lavie, 2005), and TER (Snover et al., 2006) have originated from the readily available metrics for machine translation and/or text summarization. On the contrary, the more recent metrics such as CIDEr (Vedantam et al., 2015) and SPICE (Anderson et al., 2016) are specifically developed for image caption evaluation task.

Evaluation with automatic metrics has some challenges as well. As previously analyzed in (Elliott and Keller, 2014), the existing automatic evaluation measures have proven to be inadequate in successfully mimicking the human judgements for evaluating the image descriptions. The latest eval-

Table 1: A summary of the evaluation metrics considered in this study.

Metric	Proposed to evaluate	Underlying idea
BLEU (Papineni et al., 2002)	Machine translation	n -gram precision
ROUGE (Lin, 2004)	Document summarization	n -gram recall
METEOR (Banerjee and Lavie, 2005)	Machine translation	n -gram with synonym matching
CIDEr (Vedantam et al., 2015)	Image description generation	<i>tf-idf</i> weighted n -gram similarity
SPICE (Anderson et al., 2016)	Image description generation	Scene-graph synonym matching
WMD (Kusner et al., 2015)	Document similarity	Earth Mover Distance on <i>word2vec</i>

uation results of 2015 MS COCO Challenge on image captioning has also revealed some interesting findings in line with this observation (Vinyals et al., 2016). In the challenge, the recent deep models outperform the human upper bound according to automatic measures, yet they could not beat the humans when the subjective human judgements are considered. These demonstrate that we need to better understand the drawbacks of existing automatic evaluation metrics. This motivates us to present an in-depth analysis of the current metrics employed in image description evaluation.

We first review BLEU, ROUGE, METEOR, CIDEr and SPICE metrics, and discuss their main drawbacks. In this context, we additionally describe WMD metric which has been recently proposed as a distance measure between text documents in (Kusner et al., 2015). We then investigate the performance of these automatic metrics through different experiments. We analyze how well these metrics mimic human assessments by estimating their correlations with the collected human judgements. Different from the previous related work (Elliott and Keller, 2014; Vedantam et al., 2015; Anderson et al., 2016), we perform a more accurate analysis by additionally reporting the results of Williams significance test. This further allows us to figure out the differences and/or similarities between a pair of metrics, whether any two metrics complement each other or provide similar results. We then test the ability of these metrics to distinguish certain pairs of captions from one another in reference to a ground truth caption. Next, we carry out an analysis on robustness of these metrics by analyzing how well they cope with the distractions in the descriptions (Hodosh and Hockenmaier, 2016).

2 Evaluation Metrics

A summary of the metrics investigated in our study is given in Table 1. All these metrics ex-

cept SPICE and WMD define the similarity over words or n -grams of reference and candidate descriptions by considering different formulas. On the other hand, SPICE (Anderson et al., 2016) considers a scene-graph representation of an image by encoding objects, their attributes and relations between them, and WMD leverages word embeddings to match groundtruth descriptions with generated captions.

2.1 BLEU

BLEU (Papineni et al., 2002) is one of the first metrics that have been in use for measuring similarity between two sentences. It has been initially proposed for machine translation, and defined as the geometric mean of n -gram precision scores multiplied by a brevity penalty for short sentences. In our experiments, we use the smoothed version of BLEU as described in (Lin and Och, 2004).

2.2 ROUGE

ROUGE (Lin, 2004) is initially proposed for evaluation of summarization systems, and this evaluation is done via comparing overlapping n -grams, word sequences and word pairs. In this study, we use ROUGE-L version, which basically measures the longest common subsequences between a pair of sentences. Since ROUGE metric relies highly on recall, it favors long sentences, as also noted by (Vedantam et al., 2015).

2.3 METEOR

METEOR (Banerjee and Lavie, 2005) is another machine translation metric. It is defined as the harmonic mean of precision and recall of uni-gram matches between sentences. Additionally, it makes use of synonyms and paraphrase matching. METEOR addresses several deficiencies of BLEU such as recall evaluation and the lack of explicit word matching. n -gram based measures work reasonably well when there is a significant overlap

between reference and candidate sentences; however they fail to spot semantic similarity when the common words are scarce. METEOR handles this issue to some extent using WordNet-based synonym matching, however just looking at synonyms may be too restrictive to capture overall semantic similarity.

2.4 CIDER

CIDER (Vedantam et al., 2015) is a recent metric proposed for evaluating the quality of image descriptions. It measures the consensus between candidate image description c_i and the reference sentences, which is a set $S_i = \{s_{i1}, \dots, s_{im}\}$ provided by human annotators. For calculating this metric, an initial stemming is applied and each sentence is represented with a set of 1-4 grams. Then, the co-occurrences of n -grams in the reference sentences and candidate sentence are calculated. In CIDER, similar to *tf-idf*, the n -grams that are common in all image descriptions are down-weighted. Finally, the cosine similarity between n -grams (referred as $CIDER_n$) of the candidate and the references is computed.

CIDER is designed as a specialized metric for image captioning evaluation, however, it works in a purely linguistic manner, and only extends existing metrics with *tf-idf* weighting over n -grams. This sometimes causes unimportant details of a sentence to be weighted more, resulting in a relatively ineffective caption evaluation.

2.5 SPICE

Another recently proposed metric for evaluating image caption similarity is SPICE (Anderson et al., 2016). It is based on the agreement of the scene-graph tuples (Johnson et al., 2015; Schuster et al., 2015) of the candidate sentence and all reference sentences. Scene-graph is essentially a semantic representation that parses the given sentence to semantic tokens such as object classes C , relation types R and attribute types A . Formally, a candidate caption c is parsed into a scene-graph as

$$G(c) = \langle O(c), E(c), K(c) \rangle$$

where $G(c)$ denotes the scene graph of caption c , $O(c) \subseteq C$ is the set of object mentions, $E(c) \subseteq O(c) \times R \times O(c)$ is the set of hyper-edges representing relations between objects, and $K(c) \subseteq O(c) \times A$ is the set of attributes associated with objects. Once the parsing is done, a set of tuples is




Figure 1: An example image with its Scene Graph where the parser fails to parse the candidate sentence accurately, which could result in wrong calculation of SPICE metric. See text for details.

formed by using the elements of G and their possible combinations. SPICE score is then defined as the F_1 -score based on the agreement between the candidate and reference caption tuples. For tuple matching, SPICE uses WordNet synonym matching (Pedersen et al., 2004) as in METEOR (Banerjee and Lavie, 2005). One problem is that the performance becomes quite dependent on the quality of the parsing. Figure 1 illustrates an example case of failure. Here, swimming is parsed as an object, with all its relations, and dog is parsed as an attribute.

2.6 WMD

Two captions may not share the same words or any synonyms; yet they can be semantically similar. On the contrary, two captions may include similar objects, attributes or relations yet they may not be semantically similar. Metrics that are currently in use fail to correctly identify and assess the quality of such cases. To address this issue, we propose to use a recently introduced document distance measure called Word Mover’s Distance (WMD) (Kusner et al., 2015) for evaluating image captioning approaches. WMD casts the distance between documents as an instance of Earth Mover’s Distance (EMD) (Rubner et al., 2000), where travel costs are calculated based on *word2vec* (Mikolov et al., 2013) embeddings of the words.

For WMD, text documents (in our case image captions) are first represented by their normalized bag-of-words (nBOW) vectors, accounting for all words except stopwords. More formally, each text document is represented as vectors $\mathbf{d} \in R^n$, where, $d_i = \frac{c_i}{\sum_{j=1}^n c_j}$ if a word i appears c_i times in the document. WMD incorporates semantic similarity between individual word pairs into the document similarity metric, by using the distances in

Table 2: Drawbacks of automatic evaluation metrics for image captioning. See text for details.

	Description	BLEU	METEOR	ROUGE	CIDEF	SPICE	WMD
original	a man wearing a red life jacket is sitting in a canoe on a lake	1	1	1	10	1	1
candidate	a man wearing a life jacket is in a small boat on a lake	0.45	0.28	0.67	2.19	0.40	0.19
synonyms	a guy wearing a life vest is in a small boat on a lake	0.20	0.17	0.57	0.65	0.00	0.10
redundancy	a man wearing a life jacket is in a small boat on a lake at sunset	0.45	0.28	0.66	2.01	0.36	0.18
word order	in a small boat on a lake a man is wearing a life jacket	0.26	0.26	0.38	1.32	0.40	0.19




Figure 2: An illustration of the distance calculation of WMD metric comparing two candidate captions with a reference caption.

word2vec embedding space. Specifically, the distance between word i and word j in two documents is set as the Euclidean distance between each of the corresponding *word2vec* embeddings x_i and x_j , *i.e.*, $c(i, j) = \|x_i - x_j\|_2$.

The distances between words serve as building blocks to define distances between documents, hence captions. The flow between word vectors is defined with the sparse flow matrix $\mathbf{T} \in R^{n \times n}$, with T_{ij} representing the travel amount of word i to word j . The distance between two documents is then defined with $\sum_{i,j} T_{ij} c(i, j)$, *i.e.* the minimum cumulative cost required to move all words between documents. This minimum cumulative cost is found by solving the corresponding linear optimization problem, which is cast as a special case of EMD metric (Rubner et al., 2000). An example matching result is shown in Figure 2. By using *word2vec* embeddings, semantic similarities between words are more accurately identified. In our experiments, we convert the distance scores to similarities by using a negative exponential.

2.7 Drawbacks of the metrics

In order to illustrate the drawbacks of these automatic evaluation metrics, we provide an exam-

ple case in Table 2. In this table, an original caption is given, together with the upper bound values for each metric, *i.e.* when this original caption is compared to itself. The second line includes a candidate caption that is semantically very similar to the original one and the corresponding similarity scores according to evaluation metrics. We then modify the candidate sentence slightly and observe how the metric scores are affected from these small modifications. First, we observe that all the scores decrease when some words are replaced with their synonyms. The change is especially significant for SPICE and CIDEF. In this example, failure of SPICE is likely due to incorrect parsing or the failure of synonym matching. On the other hand, failure of CIDEF is likely due to unbalanced *tf-idf* weighting. Second, we observe that the metrics are not affected much from the introduction of additional (redundant) words in the sentences. However, when the order of the words are changed, we see that BLEU, ROUGE and CIDEF scores decrease notably, due to their dependence on n -gram matching. Note that, WMD and SPICE are not influenced from the change in word order.

3 Evaluation and Discussion

3.1 Quality

A common way of assessing the performance of a new automatic image captioning metric is to analyze how well it correlates with human judgments of description quality. However, in the literature, there is no consensus on which correlation coefficient is best suited for measuring the soundness of a metric in this way. Elliott and Keller (2014) reports Spearman’s rank correlation, which measures a monotonic relation, whereas Anderson *et al.* (2016) suggests to use Pearson’s correlation, which assumes that the relation is lin-

ear, and Kendall’s correlation, which is another rank correlation measure.

The above correlation analysis is a well-established practice for automatic metric evaluation, but it is not complete in the sense that it is not meaningful to draw conclusions from it about the differences or similarities between a pair of metrics. That is, comparing the corresponding correlations relative to each other does not say much since they are both computed on the same dataset, and thus not independent. To address this issue, Graham and Baldwin (2014) have suggested to use Williams significance test (Williams, 1959), which also takes into account the degree to which the two metrics correlate with each other, and can reveal whether one metric significantly outperforms the other. The test has shown to be valuable for evaluation of document and segment-level machine translation (Graham and Baldwin, 2014; Graham et al., 2015; Graham and Liu, 2016) and summarization metrics (Graham, 2015). In this study, we extend the previous correlation-based evaluations of image captioning metrics by providing a more conclusive analysis based on Williams significance test.

Williams test (Williams, 1959) calculates the statistical significance of differences in dependent correlations, and formulated as testing whether the population correlation between X_1 and X_3 equals the population correlation between X_2 and X_3 :

$$t(n-3) = \frac{(r_{13} - r_{23})\sqrt{(n-1)(1+r_{12})}}{\sqrt{2K\left(\frac{n-1}{n-3}\right) + \frac{(r_{23}+r_{13})^2}{4}(1-r_{12})^3}} \quad (1)$$

where r_{ij} is the correlation between X_i and X_j , and n is the size of the population, with

$$K = 1 - r_{12}^2 - r_{13}^2 - r_{23}^2 + 2r_{12}r_{13}r_{23}. \quad (2)$$

To analyze statistical significance in the automatic metrics listed in Section 2, we use the publicly available FLICKR-8K (Elliott and Keller, 2014) and COMPOSITE (Aditya et al., 2015) datasets, which we describe below. We note that in our experiments, we first lowercase and tokenize the candidate and reference captions using `ptbtokenizer.py` script from MS COCO evaluation tools¹. We use the implementations of the metrics from the same evaluation kit with the ex-

¹<https://github.com/peteanderson80/coco-caption>

ception of WMD. For the WMD metric, we employ the code provided by Kusner et al. (2015)².

FLICKR-8K³ dataset contains quality judgements for 5822 candidate sentences for the images in its test set (Hodosh et al., 2013). These judgements are collected from 3 human experts and they are on a scale of [1, 4], with a score of 1 denoting a description totally unrelated to the image content, and 4 meaning a perfect description for the image. Candidate captions are all obtained from a retrieval based model, hence they are grammatically correct.

COMPOSITE⁴ dataset contains human judgements for 11,985 candidate captions for the subsets of FLICKR-8K (Hodosh et al., 2013), FLICKR-30K (Young et al., 2014) and MS COCO (Lin et al., 2014) datasets. The AMT workers were asked to judge the candidate caption for an image using two aspects: (i) *correctness*, and (ii) *thoroughness* of the candidate caption, both on a scale of [1, 5] where 1 means not relevant/less detailed and 5 denotes the candidate caption perfectly describing the image. Candidate captions were sampled from the human reference captions and the captioning models in (Aditya et al., 2015; Karpathy and Fei-Fei, 2015).

Table 3 shows Pearson’s, Spearman’s and Kendall’s correlation of the metrics with the human judgements in FLICKR-8K and COMPOSITE datasets. For FLICKR-8K, we follow the methodology in (Elliott and Keller, 2014) and compute correlations with the human expert scores. On the other hand, for COMPOSITE, we report the mean of the correlations with *correctness* and *thoroughness* scores. In terms of these correlations, while SPICE produces the highest quality comparisons in FLICKR-8K, WMD and METEOR give better results in COMPOSITE in general. However, if one further inspects the score distributions of the metrics (on FLICKR-8K dataset) shown in Figure 3, while SPICE can identify irrelevant captions remarkably well, it can not effectively distinguish bad captions from relatively better ones.

In Figure 4(a), we show Spearman’s correlation between each pair of metrics, where the metrics are ordered from highest to lowest correlation with

²<https://github.com/mkusner/wmd>

³<https://github.com/elliotttd/compareImageDescriptionMeasures>

⁴<https://imagesdg.wordpress.com/image-to-scene-description-graph>

Table 3: Correlation between automatic image captioning metrics and human judgement scores.

	FLICKR-8K			COMPOSITE		
	Pearson	Spearman	Kendall	Pearson	Spearman	Kendall
WMD	0.68	0.60	0.48	0.43	0.43	0.32
SPICE	0.69	0.64	0.56	0.40	0.42	0.34
CIDEr	0.60	0.56	0.45	0.32	0.42	0.32
METEOR	0.69	0.58	0.47	0.37	0.44	0.33
BLEU	0.59	0.44	0.35	0.34	0.38	0.28
ROUGE	0.57	0.44	0.35	0.40	0.39	0.29

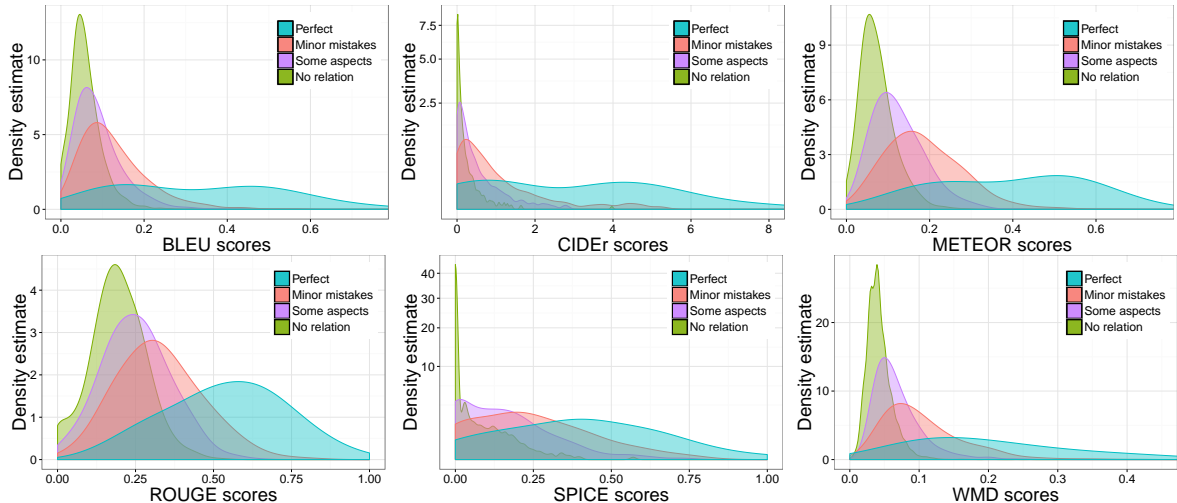


Figure 3: Score distributions of the metrics on FLICKR-8K dataset. Four different rating scales are used: 1 for no relation, 2 for minor mistakes, 3 for some true aspects and 4 for perfect match. For CIDEr and SPICE metrics, square-root transform is performed on the y -axis to better illustrate how the score distributions overlap with each other.

human judgements⁵. Overall, the pairwise correlations are generally high for both datasets. We additionally observe that the metrics which depend on similar structures are grouped together using these correlations. For example, the n -gram based metrics BLEU and ROUGE provide scores that are highly correlated with each other for FLICKR-8K. The correlations within COMPOSITE dataset are even very high for all the metrics that consider n -grams, namely BLEU, CIDEr, METEOR and ROUGE. On the other hand, the correlations of these metrics against SPICE and WMD are not that high. Moreover, the pairwise correlations between SPICE and WMD are relatively low as well. All these findings suggest that these three groups of metrics, the n -gram based metrics, the scene-graph based SPICE and the word embedding

⁵Here, we only report Spearman’s correlation since, compared to Pearson’s, it provides a more consistent ranking of the metrics across the two datasets, and is similar to Kendall’s correlation.

based WMD, can be complementary to each other.

Finally, in Figure 4(b), we provide the results of Williams significance test, which compares two different metrics with respect to their correlations against human judgements. Our results show that all the metric pairs have a significant difference in correlation with human judgement at $p < 0.05$. This reveals that the pair of metrics which has close correlation scores with human judgements (e.g. SPICE and WMD in FLICKR-8K dataset) are found to be statistically different than each other. These findings collectively support our previous conclusion that all metrics considered here can complement each other in evaluating the quality of the generated captions.

3.2 Accuracy

In this section, following the methodology introduced in (Vedantam et al., 2015), we analyze the ability of each metric to discriminate certain pair of captions from one another in reference to

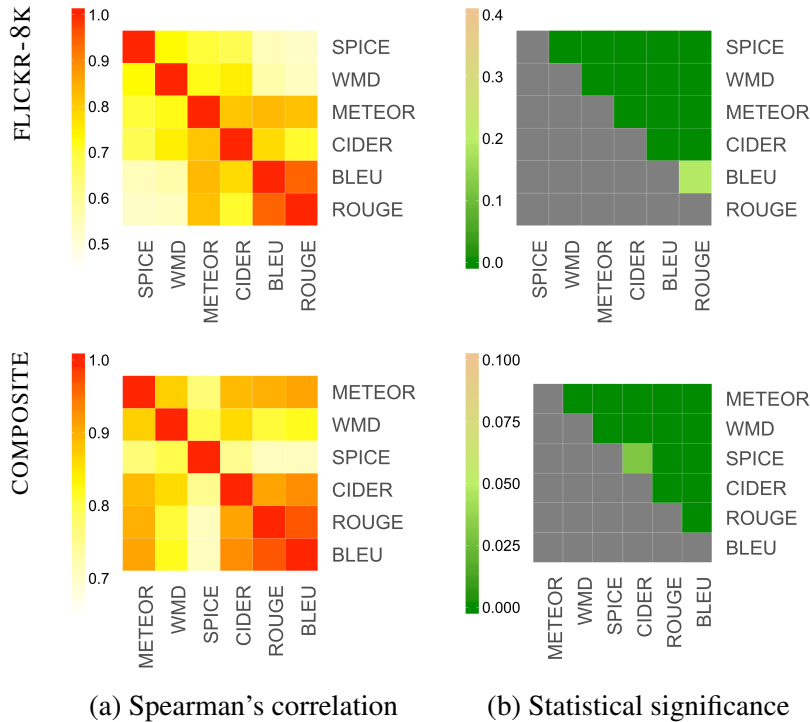


Figure 4: Significance test results for pairs of automatic metrics on FLICKR-8K and COMPOSITE datasets. (a) Spearman correlation between pairs of metrics; and (b) p -value of Williams significance tests, green cells indicate a significant win for metric in row i over metric in column j .

a groundtruth caption. We employ the human consensus scores while evaluating the accuracies. In particular, for evaluation, a triplet of descriptions, one reference and two candidate descriptions, is shown to human subjects and they are asked to determine the candidate description that is more similar to the reference. A metric is accurate if it provides a higher score to the description chosen by the human subject as being more similar to the reference caption. For this analysis, we carry out our experiments on PASCAL-50S and ABSTRACT-50S datasets⁶. We consider different kinds of pairs such as (human-human correct) HC, (human-human incorrect) HI, (human-machine) HM, and (machine-machine) MM. As the candidate sentences are generated by both humans and machines, each test scenario has a different level of difficulty.

ABSTRACT-50S (Vedantam et al., 2015) dataset is a subset of the Abstract Scenes Dataset (Zitnick and Parikh, 2013), which includes 500 images containing clipart objects in everyday scenes. Each image is annotated with 50 different descriptions. For evaluation, 48 of these 50 descriptions are used as reference descriptions and the remain-

ing 2 descriptions are employed as candidate descriptions. For 400 pairs of these descriptions, human consensus scores are available, with the first 200 are for HC and the remaining 200 are for HI.

PASCAL-50S (Vedantam et al., 2015) dataset is an extended version of the Pascal Sentences (Farhadi et al., 2010) dataset that contains 1000 images from PASCAL Object Detection challenge (Everingham et al., 2010) of 20 different object classes like person, car, horse, etc. This version includes 50 captions per image and human judgements for 4000 candidate pairs for the aforementioned binary-forced choice task, which are all collected through Amazon Mechanical Turk (AMT). For this dataset, all four different categories are available, having 1000 pairs for each category.

In Table 4, we present caption-level classification accuracy scores of automatic evaluation metrics at matching human consensus scores. On ABSTRACT-50S dataset, the CIDER metric outperforms all other metrics in both HC and HI cases. On the other hand, on PASCAL-50S dataset, the WMD metric gives the best scores in three out of four cases. Especially, it is the most accurate metric at matching human judgements on the chal-

⁶<http://vrarna91.github.io/cider>

Table 4: Description-level classification accuracies of automatic evaluation metrics.

	ABSTRACT-50S			PASCAL-50S				
	HC	HI	Avg.	HC	HI	HM	MM	Avg.
WMD	0.65	0.93	0.79	0.71	0.99	0.93	0.74	0.84
SPICE	0.62	0.89	0.76	0.66	0.98	0.85	0.72	0.81
CIDEr	0.76	0.95	0.86	0.69	0.99	0.94	0.66	0.82
METEOR	0.60	0.90	0.75	0.69	0.99	0.90	0.65	0.81
BLEU	0.69	0.89	0.79	0.67	0.97	0.94	0.60	0.80
ROUGE	0.65	0.89	0.77	0.68	0.97	0.92	0.60	0.79



Distractor Type	Gold Caption	Distracted Caption
“Replace Scene”	a man wearing a life jacket is in a small boat on a lake with a ferry in view	a man wearing a life jacket is in a small boat on takeoff with a ferry in view
“Replace Person”	a man wearing a life jacket is in a small boat on a lake with a ferry in view	a woman in a blue shirt and headscarf is in a small boat on a lake with a ferry in view
“Share Person”	a man wearing a life jacket is in a small boat on a lake with a ferry in view	a man is selecting a chair from a stack under a shady awning
“Share Scene”	a man wearing a life jacket is in a small boat on a lake with a ferry in view	a black and brown dog is playing on the ice at the edge of a lake

Figure 5: Distracted versions of an image descriptions for a sample image.

Table 5: Distraction analysis.

Case	# Instances	BLEU	METEOR	ROUGE	CIDEr	SPICE	WMD
Replace-Scene	2514	0.62	0.69	0.63	0.83	0.54	0.76
Replace-Person	5817	0.73	0.77	0.78	0.78	0.67	0.80
Share-Scene	2621	0.79	0.85	0.79	0.81	0.70	0.87
Share-Person	4596	0.78	0.85	0.78	0.83	0.67	0.88
Overall	15548	0.73	0.79	0.75	0.81	0.65	0.83

lenging MM and HC cases, which require distinguishing fine-grained differences between descriptions. On average, the performances of all the other metrics are very similar to each other.

3.3 Robustness

In this section, we evaluate the robustness of the automatic image captioning metrics. For this purpose, we employ the binary (two-alternative) forced choice task introduced in (Hodosh and Hockenmaier, 2016) to compare the existing image captioning models. For a given image, this task involves distinguishing a correct description from its slightly distracted incorrect versions. In our case, a robust image captioning metric should always choose the correct caption over the distracted ones.

In our experiments, we use the data⁷ provided

⁷http://nlp.cs.illinois.edu/HockenmaierGroup/Papers/VL2016/HodoshHockenmaier16_BinaryTasks_Data.tar

by the authors for a subset of FLICKR-30K (Hodosh et al., 2013). Specifically, we consider four different types of distractions for the image descriptions, namely 1) Replace-Scene, 2) Replace-Person, 3) Share-Scene, and 4) Share-Person, which results 15548 correct and distracted caption pairs in total. For Replace-Scene and Replace-Person tasks, the distracted descriptions were artificially constructed by replacing the main actor (first person) and the scene in the original caption by random person and scene elements, respectively. For Share-Scene and Share-Person tasks, the distracted captions were selected from the sentences from the training part of FLICKR-30K (Young et al., 2014) dataset whose actor or scene chunks share the similar main actor or scene elements with the correct description. Figure 5 presents an example image together with the original description and its distracted versions.

We compare each correct caption available for an image with the remaining correct and distracted

captions for that image by considering tested evaluation metrics, and then estimate an average accuracy score. In Table 5, we present the classification accuracies of the evaluation metrics for each distraction type. As can be seen, the WMD metric gives the best results for three out of four categories, and provides the second best result for the Replace-Scene case. Overall, METEOR and CIDER metrics seem to be also robust to these distractions. The very recently proposed SPICE metric performs the worst for this task. This is somewhat expected as it is even affected by the use of synonyms of the words as we have previously shown in Table 2.

3.4 Discussion

As the experiments on quality, accuracy and robustness tests demonstrate in Sections 3.1-3.3, existing automatic image captioning metrics all have some strengths and weaknesses due to their design choices. For example, while SPICE, METEOR and WMD give the best performances in terms of our correlation analysis against human judgments, CIDER and WMD provide the best classification scores for our accuracy experiments. Moreover, CIDER, METEOR and WMD are found to be less affected by the distractors. Overall, our analysis suggests that the recently proposed WMD document metric is also quite effective for image captioning since it has high correlations with the human scores, is much less sensitive to synonym swapping and additionally performs well at the accuracy and distraction tasks.

Our analysis also shows that the existing metrics both theoretically and empirically differ from each other with significant differences. Compared to the recent results of significance testing of machine translation and summarization metrics (Graham and Baldwin, 2014; Graham et al., 2015; Graham and Liu, 2016; Graham, 2015), our results suggest that there remains much room for improvement in developing more effective image captioning evaluation metrics. We leave this for future work, but a very naive idea would be combining different metrics into a unified metric and we simply test this idea using score combination, after normalizing the score of each metric to the range $[0, 1]$. Among all possible combinations, we find that the combination of WMD+SPICE+METEOR performs the best with a Spearman’s correlation of 0.66 for FLICKR-8K and

0.45 for COMPOSITE dataset, yielding an improvement from SPICE (0.64 and 0.42). In addition, we should add that this unified metric significantly outperforms the individual metrics according to Williams test ($p < 0.01$).

4 Conclusion

In this paper, we provide a careful evaluation of the automatic image captioning metrics, and propose to use WMD, which utilizes *word2vec* embeddings of the words to compute a semantic similarity of sentences. We highlight the drawbacks of the existing metrics, and we empirically show that they are significantly different than each other. We hope that this work motivates further research into developing better evaluation metrics, probably learning based ones, as previously studied in machine translation literature (Kotani and Yoshimi, 2010; Guzmán et al., 2015). We also observe that incorporating visual information (via Scene-graph used by SPICE) and semantic information (via WMD) is useful for the caption evaluation task, which motivates the use of multimodal embeddings (Kottur et al., 2015).

Acknowledgments

We thank the anonymous reviewers for their valuable comments. This work is supported in part by The Scientific and Technological Research Council of Turkey (TUBITAK), with award no 113E116.

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 *ECCV*.
- Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for mt evaluation with improved correlation with human judgments. In *ACL Workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, volume 29, pages 65–72.
- Raffaella Bernardi, Ruket Cakici, Desmond Elliott, Aykut Erdem, Erkut Erdem, Nazli Ikizler-Cinbis, Frank Keller, Adrian Muscat, and Barbara Plank.

2016. Automatic description generation from images: A survey of models, datasets, and evaluation measures. *JAIR*, 55:409–442.
- Desmond Elliott and Frank Keller. 2013. Image description using visual dependency representations. In *EMNLP*, volume 13, pages 1292–1302.
- Desmond Elliott and Frank Keller. 2014. Comparing automatic evaluation measures for image description. In *ACL*, pages 452–457.
- Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. 2010. The Pascal visual object classes (VOC) challenge. *IJCV*, 88(2):303–338.
- Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. 2010. Every picture tells a story: Generating sentences from images. In *ECCV*, pages 15–29.
- Yvette Graham and Timothy Baldwin. 2014. Testing for significance of increased correlation with human judgment. In *EMNLP*, pages 172–176.
- Yvette Graham and Qun Liu. 2016. Achieving accurate conclusions in evaluation of automatic machine translation metrics. In *NAACL-HLT*.
- Yvette Graham, Nitika Mathur, and Timothy Baldwin. 2015. Accurate evaluation of segment-level machine translation metrics. In *NAACL-HLT*.
- Yvette Graham. 2015. Re-evaluating automatic summarization with BLEU and 192 shades of ROUGE. In *EMNLP*, pages 128–137.
- Francisco Guzmán, Shafiq Joty, Lluís Márquez, and Preslav Nakov. 2015. Pairwise neural machine translation evaluation. In *ACL-IJNLP*, pages 805–814.
- Micah Hodosh and Julia Hockenmaier. 2016. Focused evaluation for image description with binary forced-choice tasks. In *VL*.
- Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013. Framing image description as a ranking task: Data, models and evaluation metrics. *JAIR*, 47:853–899.
- Xu Jia, Efstratios Gavves, Basura Fernando, and Tinne Tuytelaars. 2015. Guiding long-short term memory for image caption generation. *arXiv preprint arXiv:1509.04942*.
- Justin Johnson, Ranjay Krishna, Michael Stark, Li-Jia Li, David A Shamma, Michael S Bernstein, and Li Fei-Fei. 2015. Image retrieval using scene graphs. In *CVPR*, pages 3668–3678.
- Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, pages 3128–3137.
- Katsunori Kotani and Takehiko Yoshimi. 2010. A machine learning-based evaluation method for machine translation. In *SETN*, pages 351–356.
- Satwik Kottur, Ramakrishna Vedantam, José MF Moura, and Devi Parikh. 2015. Visual Word2Vec (vis-w2v): Learning visually grounded word embeddings using abstract scenes. *arXiv preprint arXiv:1511.07067*.
- Girish Kulkarni, Visruth Premraj, Vicente Ordonez, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C Berg, and Tamara L Berg. 2013. Babytalk: Understanding and generating simple image descriptions. *IEEE Trans. Pattern Anal. Mach. Intel.*, 35(12):2891–2903.
- Matt J. Kusner, Yu Sun, Nicholas I. Kolkin, and Kilian Q. Weinberger. 2015. From word embeddings to document distances. In *ICML*, pages 957–966.
- Siming Li, Girish Kulkarni, Tamara L. Berg, Alexander C. Berg, and Yejin Choi. 2011. Composing simple image descriptions using web-scale n-grams. In *CoNLL*, pages 220–228.
- Chin-Yew Lin and Franz Josef Och. 2004. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *ACL*, pages 605–612.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *ECCV*, pages 740–755.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text summarization branches out: Proceedings of the ACL-04 workshop*, volume 8.
- T. Mikolov, K. Chen, G. Corrado, and J. Dean. 2013. Distributed representations of words and phrases and their compositionality. In *NIPS*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *ACL*, pages 311–318.
- Ted Pedersen, Siddharth Patwardhan, and Jason Michellizzi. 2004. Wordnet::similarity: measuring the relatedness of concepts. In *Demonstration papers at HLT-NAACL 2004*, pages 38–41.
- Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. 2000. The earth mover’s distance as a metric for image retrieval. *IJCV*, 40(2):99–121.
- Sebastian Schuster, Ranjay Krishna, Angel Chang, Li Fei-Fei, and Christopher D Manning. 2015. Generating semantically precise scene graphs from textual descriptions for improved image retrieval. In *VL*, pages 70–80.

- Matthew Snover, Bonnie Dorr, and Richard Schwartz. 2006. A study of translation edit rate with targeted human annotation. In *AMTA*.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *CVPR*, pages 4566–4575.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *CVPR*, pages 3156–3164.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2016. Show and Tell: Lessons learned from the 2015 MSCOCO image captioning challenge. *arXiv preprint arXiv:1609.06647*.
- Evan J. Williams. 1959. *Regression analysis*, volume 14. Wiley New York.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhutdinov, Richard S. Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. *arXiv preprint arXiv:1502.03044*.
- 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. *TACL*, 2:67–78.
- C Lawrence Zitnick and Devi Parikh. 2013. Bringing semantics into focus using visual abstraction. In *CVPR*, pages 3009–3016.