

# Learning to Evaluate Image Captioning

Yin Cui<sup>1,2</sup> Guandao Yang<sup>1</sup> Andreas Veit<sup>1,2</sup> Xun Huang<sup>1,2</sup> Serge Belongie<sup>1,2</sup>  
<sup>1</sup>Department of Computer Science, Cornell University <sup>2</sup>Cornell Tech

## Abstract

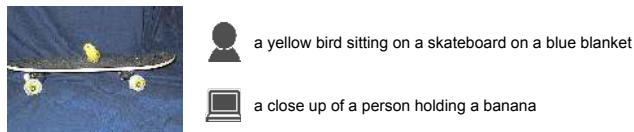
Evaluation metrics for image captioning face two challenges. Firstly, commonly used metrics such as CIDEr, METEOR, ROUGE and BLEU often do not correlate well with human judgments. Secondly, each metric has well known blind spots to pathological caption constructions, and rule-based metrics lack provisions to repair such blind spots once identified. For example, the newly proposed SPICE correlates well with human judgments, but fails to capture the syntactic structure of a sentence. To address these two challenges, we propose a novel learning based discriminative evaluation metric that is directly trained to distinguish between human and machine-generated captions. In addition, we further propose a data augmentation scheme to explicitly incorporate pathological transformations as negative examples during training. The proposed metric is evaluated with three kinds of robustness tests and its correlation with human judgments. Extensive experiments show that the proposed data augmentation scheme not only makes our metric more robust toward several pathological transformations, but also improves its correlation with human judgments. Our metric outperforms other metrics on both caption level human correlation in Flickr 8k and system level human correlation in COCO. The proposed approach could be served as a learning based evaluation metric that is complementary to existing rule-based metrics.

## 1. Introduction


Learning to automatically generate captions to summarize the content of an image is considered as a crucial task in Computer Vision. The evaluation of image captioning models is generally performed using metrics such as BLEU [27], METEOR [20], ROUGE [23] or CIDEr [31], all of which mainly measure the word overlap between generated and reference captions. The recently proposed SPICE [3] measures the similarity of scene graphs constructed from the candidate and reference sentence, and shows better correlation with human judgments.

These commonly used evaluation metrics face two challenges. Firstly, many metrics fail to correlate well with hu-

### Labeled training examples



### Learned critique



### Caption Evaluation

	Caption	Score
	a cat is watching a television on a television	0.1
	a cat is sitting on top of a television	0.15
	a cat watching birds flying on a sony tv screen	0.95
	a cat that is looking at a television	0.5

Figure 1. An overview of our proposed captioning evaluation metric. From a set of images and corresponding human written and machine generated captions, we train a model to discriminate between human and generated captions. The model comprises three major components: a CNN to compute image representations, an RNN with LSTM cells to encode the caption, and a binary classifier as the critique. After training, the learned critique can be used as a metric to evaluate the quality of candidate captions with respect to the context (*i.e.*, the image and reference human captions).

man judgments. Metrics based on measuring word overlap between candidate and reference captions find it difficult to capture semantic meaning of a sentence, therefore often lead to bad correlation with human judgments. Secondly, each evaluation metric has its well-known blind spot, and rule-based metrics are often inflexible to be responsive to new pathological cases. For example, SPICE is sensitive to the semantic meaning of a caption but tends to ignore its syntactic quality. Liu *et al.* [25] shows that SPICE prefers to give high score to long sentences with repeating clauses. It's not easy to let SPICE take such pathological cases into account. Since it's difficult to completely avoid such blind spots, a good evaluation metric for image captioning should

be flexible enough to adapt to pathological cases once identified, while correlating well with human judgments.

To address the aforementioned two challenges, we propose a metric that directly discriminates between human and machine generated captions while being able to flexibly adapt to pathological cases of our interests. Since real human judgment is impractical to obtain at scale, our proposed learning based metric is trained to perform like a human critique, as illustrated in Fig. 1. We use a state-of-the-art CNN architecture to capture high-level image representations, and a RNN with LSTM cells to encode captions. To design the learned critique, we follow insights from the COCO Captioning Challenge in 2015 [1, 7], in which a large-scale human judgment experiment was performed. In particular, our critique is a binary classifier that makes a Turing Test type judgment in which it differentiates between human-written and machine-generated captions.

In order to capture targeted pathological cases, we propose to incorporate these pathological sentences as negative training examples. To systematically create such pathological sentences, we define several transformations to generate unnatural sentences that might get high scores in an evaluation metric. Our proposed data augmentation (Sec. 3.3) scheme uses these transformations to generate large number of negative examples, which guide our metric to explore a variety of possible sentence constructions that are rare to be found in real world data. Further, we propose a systematic approach to measure the robustness of an evaluation metric to a given pathological transformation (Sec. 3.4). Extensive experiments (Sec. 4) verify the effectiveness and robustness of our proposed evaluation metric and demonstrate better correlation with human judgments on COCO and Flickr 8k, compared with commonly-used image captioning metrics.

Our key contributions can be summarized as follows:

- We propose a novel learned based captioning evaluation metric that directly captures human judgments while being flexible to targeted pathological cases.
- We demonstrate key factors for how to successfully train a good captioning evaluation metric.
- We conduct comprehensive studies that demonstrates the effectiveness of the proposed metric, in particular its correlation to human judgment and robustness toward pathological transformations.

## 2. Related Work

**Captioning evaluation.** Despite recent interests, image captioning is notoriously difficult to evaluate due to the inherent ambiguity. Human evaluation scores are reliable but costly to obtain. Thus, current image captioning models are usually evaluated with automatic metrics instead of human judgments. Commonly used evaluation metrics BLEU [27], METEOR [20], ROUGE [23] and CIDEr [31] are mostly based on n-gram overlap and tend to be insensitive to se-

mantic information. Anderson *et al.* recently proposed the SPICE [3] that is based on scene graph similarity. Although SPICE obtains significantly higher correlation with human judgments, it encounters difficulties with repetitive sentences, as pointed out in [25]. It is worth noting that all above mentioned metrics rely solely on similarity between candidate and reference captions, without taking the image into consideration. Our proposed metric, on the other hand, takes image feature as input. While all the previous metrics are rule-based, our proposed metric learns to score candidate captions by training to distinguish positive and negative examples. Moreover, our proposed training scheme could flexibly take new pathological cases into account, yet traditional metrics find it hard to adapt.

**Adversarial training and evaluation.** Generative Adversarial Networks (GANs) [12] have been recently applied to generate image captions [8, 22, 6, 30]. Although GANs could provide discriminators to tell apart human and machine generated captions, they differ from our works as our discriminator focuses on evaluation instead of generation. All existing adversarial evaluation approaches define the generator performance to be inversely proportional to the classification performance of the discriminator, motivated by the intuition that a good generator should produce outputs that are hard for the discriminator to distinguish from real data. The specific configurations differ among the approaches. Im *et al.* [16] propose to train a pair of GANs and interchange their opponents during testing. Iowe *et al.* [26] attempt to train a single discriminator on a large corpus of dialogue responses generated by different dialogue systems. Other approaches [17, 4, 21] train one discriminator separately for each model. Different from implicitly generated negative examples by a generator in these work, we incorporate explicitly defined pathological transformations to generate negative examples. Moreover, none of the above literature has verified the effectiveness of their metrics by the correlation with human judgments.

## 3. Discriminative Evaluation

A caption is considered of high quality if it is judged well by humans. In particular, the quality of a generated caption is measured by how successful it can fool a critique into believing it is written by human.

### 3.1. Evaluation Metric

The proposed evaluation metric follows the general setup of a Turing Test. First, we train an automatic critique to distinguish generated captions from human-written ones. We then score candidate captions by how successful they are in fooling the critique.

Formally, given a critique parametrized by  $\Theta$ , a reference image  $i$ , and a generated caption  $\hat{c}$ , the score is defined as the probability for the caption of being human-written, as

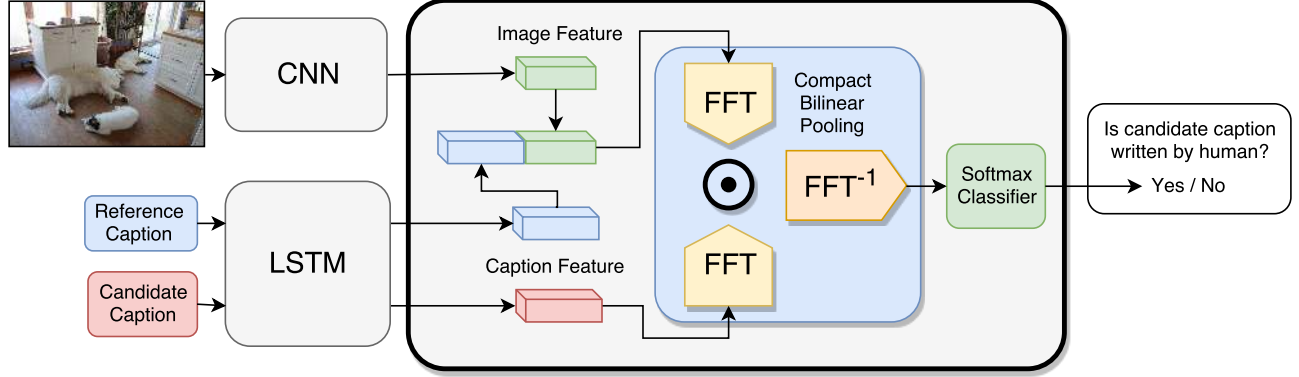


Figure 2. The model architecture of the proposed learned critique with Compact Bilinear Pooling. We use a deep residual network and an LSTM to encode the reference image and human caption into context vector. The identical LSTM is applied to get the encoding of a candidate caption. The context feature and the feature extracted from the candidate caption are combined by compact bilinear pooling. The classifier is supervised to perform a Turing Test by recognizing whether a candidate caption is human written or machine generated.

assigned by the critique:

$$\text{score}_{\Theta}(\hat{c}, i) = P(\hat{c} \text{ is human written} \mid i, \Theta) \quad (1)$$

The score is conditioned on the reference image, because the task of the evaluation is not simply to decide whether a given sentence is written by a human or machine generated, but also to evaluate whether it accurately captures the image content and focuses on the important aspects of the image.

More generally, the reference image represents the context in which the generated caption is evaluated. To provide further information about the relevance and salience of the image content, a reference caption can additionally be supplied to the context. Let  $\mathcal{C}(i)$  denotes the context of image  $i$ , then reference caption  $c$  could be included as part of context, *i.e.*,  $c \in \mathcal{C}(i)$ . The score with context becomes

$$\text{score}_{\Theta}(\hat{c}, i) = P(\hat{c} \text{ is human written} \mid \mathcal{C}(i), \Theta) \quad (2)$$

### 3.2. Model Architecture

The proposed model can be generally described in two parts. In the first part, the context information including the image and reference caption are encoded as feature vectors. These two feature vectors are then concatenated as a single context vector. In the second part, the candidate caption is encoded into a vector, in the same way as the reference caption. We then fed it into a binary classifier, together with the context vector. Fig. 2 gives an overview of the model.

To encode the image  $i$  as a feature vector  $\mathbf{i}$ , we use a ResNet [13] pre-trained on ImageNet [9] with fixed weights. The reference caption  $c$  as well as the candidate caption  $\hat{c}$  are encoded as feature vectors  $\mathbf{c}$  and  $\hat{\mathbf{c}}$  using an LSTM-based [14] sentence encoder. To form the input of the LSTM, each word is represented as a  $d$ -dimensional word embedding vector  $\mathbf{x} \in \mathbb{R}^d$  which is initialized from GloVe [28]. The LSTMs used to encode the two captions

share the same weights. The weights of the initial word embedding as well as of the LSTM are updated during training.

Once the encoded feature vectors are computed, they are combined into a single vector. In our experiments, we use two different ways to combine these features; both methods provide comparable results. The first method simply concatenates the vectors followed by a MLP:

$$\mathbf{v} = \text{ReLU}(W \cdot \text{concat}([\mathbf{i}, \mathbf{c}, \hat{\mathbf{c}}]) + b) \quad (3)$$

where  $\text{ReLU}(x) = \max(x, 0)$ . For the second method, we first concatenate the context information as  $\text{concat}([\mathbf{i}, \mathbf{c}])$  and subsequently combine it with the candidate caption using Compact Bilinear Pooling (CBP) [11], which has been demonstrated in [10] to be very effective in combining heterogeneous information of image and text. CBP uses Count Sketch [5, 29] to approximate the outer product between two vectors in a lower dimensional space. This results in a feature vector  $\mathbf{v}$  that captures 2<sup>nd</sup> order feature interactions compactly as represented by:

$$\mathbf{v} = \Phi(\text{concat}([\mathbf{i}, \mathbf{c}])) \otimes \Phi(\hat{\mathbf{c}}) \quad (4)$$

where  $\Phi(\cdot)$  represents Count Sketch and  $\otimes$  is the circular convolution. In practice, circular convolution is usually calculated in frequency domain via Fast Fourier Transform (FFT) and its inverse ( $\text{FFT}^{-1}$ ).

The feature combination is followed by a 2-way softmax classifier representing the class probabilities of being human written or machine generated. Finally, the classifier is trained using the cross-entropy loss function  $\mathcal{H}(\cdot, \cdot)$ :

$$\mathcal{L} = \frac{1}{N} \sum_{n=1}^N \mathcal{H}(p_n, q_n) \quad (5)$$

where  $N$  is the number of training examples,  $p$  is the output of the softmax classifier and  $q$  is a one-hot vector indicating




Figure 3. Example ground truth captions before and after transformations. Robustness of our learned metrics is evaluated on human captions from similar images ( $\mathcal{T}_{RC}$ ) as well as with random ( $\mathcal{T}_{RW}$ ) and permuted ( $\mathcal{T}_{WP}$ ) words.

the ground truth of whether a candidate caption is indeed human written or machine generated.

By assigning a loss function that directly captures human judgment, the learned metric is capable of measuring the objective of the image captioning task. During inference, the probability from the softmax classifier of being the human written class is used to score candidate captions.

### 3.3. Data sampling and augmentation

We would like to use data augmentation to incorporate pathological cases as negative examples during training. We define several transformations of the training data to generate a large amount of pathological sentences. Formally, a transformation  $\mathcal{T}$  takes an image-caption dataset and generates a new one:

$$\mathcal{T}(\{(c, i) \in \mathcal{D}\}; \gamma) = \{(c'_1, i'_1), \dots, (c'_n, i'_n)\} \quad (6)$$

where  $i, i'_i$  are images,  $c, c'_i$  are captions,  $\mathcal{D}$  is a list of caption-image tuples representing the original dataset, and  $\gamma$  is a hyper-parameter that controls the strength of the transformation. Specifically, we define following three transformations to generate pathological image-captions pairs:

**Random Captions(RC).** To ensure our metric pays attention to the image content, we randomly sample human written captions from other images in the training set:

$$\mathcal{T}_{RC}(\mathcal{D}; \gamma) = \{(c', i) | (c, i), (c', i') \in \mathcal{D}, i' \in \mathcal{N}_\gamma(i)\} \quad (7)$$

where  $\mathcal{N}_\gamma(i)$  represents the set of images that are top  $\gamma$  percent nearest neighbors to image  $i$ .

**Word Permutation(WP).** To make sure that our metric pays attention to sentence structure, we randomly permute at least 2 words in the reference caption:

$$\mathcal{T}_{WP}(\mathcal{D}; \gamma) = \{(c', i) | (c, i) \in \mathcal{D}, c' \in \mathcal{P}_\gamma(c) \setminus \{c\}\} \quad (8)$$

where  $\mathcal{P}_\gamma(c)$  represents all sentences generated by permuting  $\gamma$  percent of words in caption  $c$ .

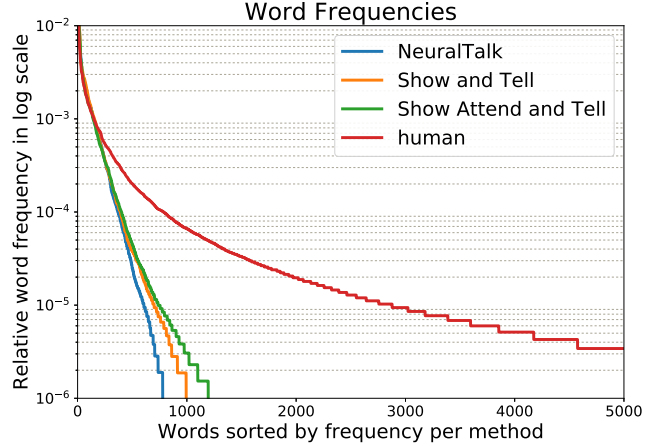


Figure 4. Relative word frequency (in log scale) in the captions generated by “NeuralTalk” [18], “Show and Tell” [32], “Show, Attend and Tell” [33], and human captions. Machine generated captions have drastically different word frequency distributions from human written captions, as human captions tend to contain much more infrequent words. As a result, a discriminator could simply detect the rare words and achieve low classification loss.

**Random Word (RW).** To explore rare words we replace from 2 to all words of the reference caption with random words from the vocabulary:

$$\mathcal{T}_{RW}(\mathcal{D}; \gamma) = \{(c', i) | (c, i) \in \mathcal{D}, c' \in \mathcal{W}_\gamma(c) \setminus \{c\}\} \quad (9)$$

where  $\mathcal{W}_\gamma(c)$  represents all sentences generated by randomly replacing  $\gamma$  percent words from caption  $c$ .

Note that all the  $\gamma$ 's are specifically defined to be a percentage.  $\gamma\% = 0$  denotes the original caption without transformation, while  $\gamma\% = 1$  provides the strongest possible transformations. Fig. 3 shows example captions before and after these transformations.

The need for data augmentation can be further illustrated by observing the word frequencies. Fig. 4 shows the relative word frequency in the captions generated by three popular captioning models as well as the frequency in human captions. Apparently, a discriminator can easily tell human and generated captions apart by simply looking at what words are used. In fact, a simple critique only trained on human written and machine-generated captions tends to believe that a sequence of random words is written by a human, simply because it contains many rare words. To address this problem, our augmented data also includes captions generated using Monte Carlo Sampling, which contains a much higher variety of words.

### 3.4. Performance Evaluation

A learned critique should be capable of correctly distinguishing human written captions from machine generated ones. Therefore, the objective of the critique is to assign

scores close to 0 to generated captions and scores close to 1 to human captions. In light of this, we define the performance of a critique as how close it gets to the ideal objectives, which is either the score assigned to a human caption or one minus the score assigned to a generated caption:

$$s(\Theta, (\hat{c}, i)) = \begin{cases} 1 - \text{score}_{\Theta}(\hat{c}, i), & \text{if } \hat{c} \text{ is generated} \\ \text{score}_{\Theta}(\hat{c}, i), & \text{otherwise} \end{cases}$$

where  $\Theta$  represents the critique,  $\hat{c}$  is the candidate caption, and  $i$  is the image  $\hat{c}$  summarizes. The performance of a model is then defined as the averaged performance on all the image-caption pairs in a test or validation set:

$$s(\Theta, \mathcal{D}) = \frac{1}{|\mathcal{D}|} \sum_{(\hat{c}, i) \in \mathcal{D}} s(\Theta, (\hat{c}, i)) \quad (10)$$

where  $\mathcal{D}$  is the set of all image-caption pairs in a held-out validation or test set.

Given a pathological transformation  $\mathcal{T}$  and  $\gamma$ , we could compute the average score of a metric  $\Theta$  on the transformed validation set  $\mathcal{T}(\mathcal{D}, \gamma)$ , *i.e.*  $s(\Theta, \mathcal{T}(\mathcal{D}, \gamma))$ . We define the robustness score with respect to transformation  $\mathcal{T}$  as the Area-Under-Curve (AUC) of  $s(\Theta, \mathcal{T}(\mathcal{D}, \gamma))$  by varying all possible  $\gamma$ :

$$R(\Theta, \mathcal{T}) = \int s(\Theta, \mathcal{T}(\mathcal{D}, \gamma)) d\gamma \quad (11)$$

We expect a robust evaluation metric to give low scores to the image-caption pairs generated by the pathological transformations. To compare metrics with different scales, we normalize the scores given by each metric such that the ground truth human caption receives a score of 1. Detailed experiments are presented in Sec. 4.2 and Sec. 4.3.

### 3.5. Using the Learned Metrics

To use the learned metrics in practice, one needs to first fix both the model architecture of the discriminator and all the hyper-parameters of the training process. When evaluating a captioning model, we need the generated captions of the model for a set of images (*i.e.*, validation or test set of a image captioning dataset). We then split the results into two folds. The discriminative metric is trained with image-caption pairs in first fold as training data, together with ground truth captions written by human. Then we use the trained metric to score the image-caption pairs on the other fold. Similarly, we score all the image-caption pairs in the first fold using a metric trained from the second fold. Once we get all the image-caption pairs scored in the dataset, the average score will be used as the evaluation of the captioning model. One could reduce the variance of the evaluation score by training the metric multiple times and use the averaged evaluation score across all the runs.

## 4. Experiments

### 4.1. Experiment setup

**Data.** We use the COCO dataset [24] to evaluate the performance of our proposed metric. To test the capability (Sec. 4.2) and robustness (Sec. 4.3) of the proposed models, we use the data split from [18], which re-splits the original COCO dataset into a training set with 113,287 images, a validation and a test set, each contains 5,000 images. Each image is annotated by roughly 5 human annotators. We use the validation set for parameter tuning. For the system level human correlation study (Sec. 4.5), we use 12 submission entries from the 2015 COCO Captioning Challenge on the COCO validation set <sup>1</sup>.

The caption level human correlation study (Sec. 4.4) uses human annotations in Flickr 8k dataset [15]. Flickr 8k collects two sets of human annotations, each on a different set of image caption pairs. Among these image-caption pairs, candidate captions are sampled from human captions in the dataset. In the first set of human annotation (Expert Annotation), human experts are asked to rate the image-caption pairs with scores ranging from 1: The selected caption is unrelated to the image to 4: The selected caption describes the image without any errors. The second set of annotation (Crowd Flower Annotation) is collected by asking human raters to decide whether a caption describes the corresponding image or not.

**Image Captioning Models.** We use publicly available implementations of “NeuralTalk” (NT) [18], “Show and Tell” (ST) [32], “Show, Attend and Tell” (SAT) [33] as image captioning models to train and evaluate our metric.

**Implementation Details.** Our image features are extracted from a Deep Residual Network with 152 layers (ResNet-152) [13] pre-trained on ImageNet. We follow the preprocessing from [32, 33, 18] to prepare vocabulary on COCO dataset. We fix the step size of the LSTM to be 15, padding shorter sentences with a special token while cutting longer ones to 15 words. All words are represented as 300-dimensional vectors initialized from GloVe [28]. We use a batch size of 100 and sample an equal number of positive and negative examples in each batch. Linear projection is used to reduce the dimension of image feature to match that of caption features. For Compact Bilinear Pooling, we use the feature dimension of 8192 as suggested in [11]. We use 1 LSTM layer with a hidden dimension of 512 in all experiments unless otherwise stated. All the model are trained using Adam [19] optimizer for 30 epochs with an initial learning rate of  $10^{-3}$ . We decay the learning rate by a factor of 0.9 after every epoch. Our code (in Tensorflow [2]) is available at: <https://github.com/richardaecn/cvpr18-caption-eval>.

<sup>1</sup>Among 15 participating teams, 3 didn’t provide submissions on validation set. Thus, we use submission entries from the remaining 12 teams.

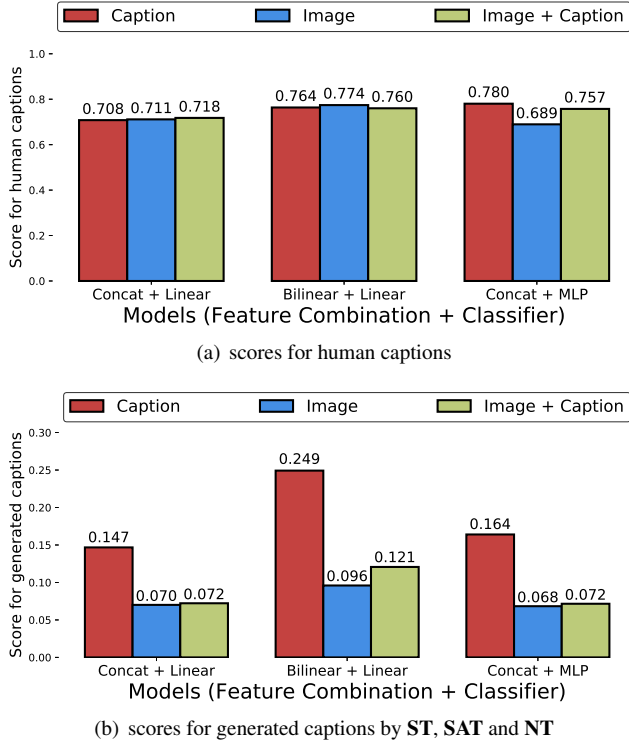


Figure 5. Top: Average score of human captions from the validation set. Bottom: Average score of generated captions. Color of the bar indicates what context information is used for the critique. The horizontal axis represents three different strategies for the combination of the features from candidate caption with the context as well as the used classifier (Concat + Linear: concatenation and linear classifier; Bilinear + Linear: compact bilinear pooling and linear classifier; Concat + MLP: concatenation and MLP with one hidden layer of size 512).

## 4.2. Capability

To measure the capability of our metric to differentiate between human and generated captions, we train variants of models using generated captions from **ST**, **SAT** and **NT**, together with human captions from the training set. Fig. 5(a) and Fig. 5(b) show the average score on the validation set for human and generated captions respectively. The results show that all models give much higher scores to human captions than machine generated captions, indicating that they are able to differentiate human-written captions from the machine-generated ones.

With respect to the choice of context, we observe that including image features into the context clearly improves performance. Also, adding a reference caption does not lead to a significant improvement over only using image features. This indicates that the image itself provides enough contextual information for the critique to successfully discriminate between human and machine generated captions. The reason that none of the commonly used metrics in-

cludes images as context is likely due to the difficulties of capturing image-text similarity. Our metric circumvents this issue by implicitly learning the image-text relationship directly from the data.

It is worth noting that achieving high model performance in terms of discrimination between human and generated captions does not necessarily imply that the learned metric is good. In fact, we observe that a critique trained without data augmentation can achieve even higher discrimination performance. Such critique, however, also gives high scores to human written captions from other images, indicating that the classification problem is essentially reduced to putting captions into categories of human and non-human written without considering the context image. If trained with the proposed data sampling and augmentation technique, the critique learns to pay attention to image context.

## 4.3. Robustness

To evaluate whether the proposed metric can capture pathological image-caption pairs, we conduct robustness studies as described in Sec. 3.4 on the three pathological transformations defined in Sec. 3.3. The robustness comparisons are illustrated in Fig. 6. In the first row we compare different variants of the proposed metric. The results illustrate that, although achieving high discrimination performance, a metric learned without data sampling or augmentation also gives high scores to human captions from other images ( $\mathcal{T}_{RC}$ ), with random words ( $\mathcal{T}_{RW}$ ), or word permutations ( $\mathcal{T}_{WP}$ ). This indicates that the model tends to focus on an overall human vs. non-human classification without considering contextual information in the image or the syntactic structure of the candidate sentence.

Further, even with data augmentation, a linear model with concatenated context and candidate caption features gives high scores to human captions from other images, possibly because there is no sufficient interaction between the context and candidate caption features. Non-linear interactions such as Compact Bilinear Pooling or a non-linear classifier with hidden layers solve this limitation. The non-linear model in Fig. 6 refers to a model with concatenated context and candidate features followed by a nonlinear classifier. Compact bilinear pooling (not shown in the figure for clarity of visualization) achieves similar results.

In the second row of Fig. 6 we compare our metric with other commonly used image captioning metrics. The proposed metric outperforms all others with respect to random ( $\mathcal{T}_{RW}$ ) as well as permuted ( $\mathcal{T}_{WP}$ ) words and is reasonably robust to human captions from similar images ( $\mathcal{T}_{RC}$ ). Further, we observe that the recently proposed metrics CIDEr and SPICE perform well for human captions from similar images, but fall behind with respect to sentence structure. This could be caused by their increased focus on informative and scene specific words.

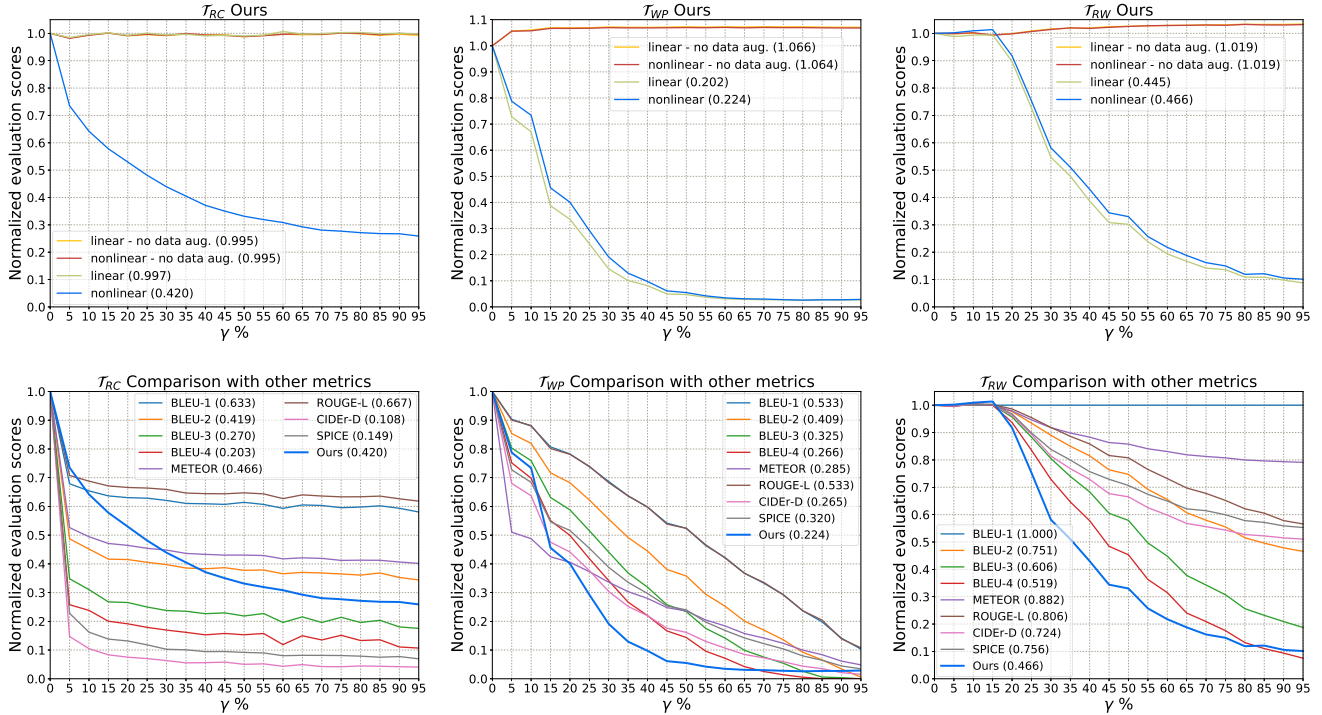


Figure 6. Normalized evaluation score for transformations  $\mathcal{T}_{RC}$  (human caption from other images),  $\mathcal{T}_{WP}$  (random permutation of words), and  $\mathcal{T}_{RW}$  (words replaced by random words) with different amount of transformation ( $\gamma\%$ ). When  $\gamma\% = 0\%$ , the original dataset is kept unchanged; when  $\gamma\% = 100\%$ , maximum amount of transformation is applied to the dataset. The first row shows results of our metrics using either linear or non-linear model trained with or without data augmentation. The second row compares our non-linear model trained with data augmentation to other metrics. The score after each metric shows the robustness score defined in Sec. 3.4, *i.e.*, the Area Under Curve (AUC). The lower the score the more robust the metric is.

#### 4.4. Caption Level Human Correlation

We use both the Expert Annotations and the Crowd Flower Annotations from Flickr 8k dataset [15] to compute caption level correlation with human judgments. We follow the procedure in SPICE paper [3] to compute the Kendall’s  $\tau$  rank correlation in the Expert Annotations. The  $\tau$  correlation for the Crowd Flower Annotation is computed between scores generated by the evaluation metric and percentage of raters who think that the caption describes the image with possibly minor mistakes. During training, all negative samples are generated by transformation  $\mathcal{T}_{RC}$ , *i.e.*, human caption from random image.

The results in Table 1 show that our metrics achieve the best caption level correlation in both Expert Annotations and Crowd Flower Annotations. Note that the Crowd Flower Annotations use a binary rating setup, while the setup from the Expert Annotations makes a finer-grained ratings. Despite the fact that our model is trained on a simpler binary objective, it still correlates well with human judgments from the Expert Annotations. Note that we do not use any human annotations during the training, since all of our negative examples could be generated automatically.

#### 4.5. System Level Human Correlation

We compare our metric with others on the Pearson’s  $\rho$  correlation between all common metrics and human judgments collected in the 2015 COCO Captioning Challenge [1]. In particular, we use two human judgment M1: Percentage of captions that are evaluated as better or equal to human caption and M2: Percentage of captions that pass the Turing Test. We don’t use M3: correctness, M4: detailness and M5: salience, as they are not used to rank image captioning models, but are intended for an ablation study to understand which aspects make captions good.

Since we don’t have access to the COCO test set annotations, where the human judgments are collected on, we perform our experiments on the COCO validation set. There are 15 teams participated in the 2015 COCO captioning challenge and we use 12 of them that submitted results on the validation set. We assume the human judgment on the validation set is sufficiently similar to the judgment on the test set. We don’t use any additional training data besides the submission files on the validation set and the data augmentation described in Sec. 3.3. To get evaluation scores on the whole validation set, we split the set in two halves and,

	Expert Annotations	Crowd Flower
BLEU-1	0.191*	0.206
BLEU-2	0.212	0.212
BLEU-3	0.209	0.204
BLEU-4	0.206*	0.202
METEOR	0.308*	0.242
ROUGE-L	0.218*	0.217
CIDEr	0.289*	0.264
SPICE	0.456	0.252
<b>Ours</b>	<b>0.466</b>	<b>0.295</b>
Inter-human	0.736	-

**Expert Annotations:** experts score image-caption pairs from 1 to 4; 1 means caption doesn't describe the image.

**Crowd Flower:** human raters mark 1 if the candidate caption describes the image, and mark 0 if not.

Table 1. Caption level Kendall's  $\tau$  correlation between Flickr 8K [15]'s human annotations and evaluation metrics' scores. Our reported scores with \* differ from the ones reported in SPICE [3].

	M1		M2	
	$\rho$	$p$ -value	$\rho$	$p$ -value
BLEU-1	0.124	(0.687)	0.135	(0.660)
BLEU-2	0.037	(0.903)	0.048	(0.877)
BLEU-3	0.004	(0.990)	0.016	(0.959)
BLEU-4	-0.019	(0.951)	-0.005	(0.987)
METEOR	0.606	(0.028)	0.594	(0.032)
ROUGE-L	0.090	(0.769)	0.096	(0.754)
CIDEr	0.438	(0.134)	0.440	(0.133)
SPICE	0.759	(0.003)	0.750	(0.003)
<b>Ours (no DA)</b>	<b>0.821</b>	(0.000)	<b>0.807</b>	(0.000)
<b>Ours</b>	<b>0.939</b>	(0.000)	<b>0.949</b>	(0.000)

**M1:** Percentage of captions that are evaluated as better or equal to human caption.

**M2:** Percentage of captions that pass the Turing Test.

Table 2. Pearson's  $\rho$  correlation between human judgments and evaluation metrics. The human correlation of our proposed metric surpasses all other metrics by large margins. Scores reported in SPICE [3] were calculated on the COCO test set for all 15 teams, whereas ours were from 12 teams on the COCO validation set.

for each submission, train our critique on each split and get scores (probability of being human written) on the other.

The results in Table 2 show that our learned metric surpasses all other metrics including the recently proposed SPICE [3] by large margins, especially trained with data augmentation. This indicates that aligning the objective with human judgments and using data augmentation yield a better evaluation metric. Fig. 7 illustrates our metric compared with human judgment - M1 on COCO validation set. Our metric aligns well with human judgment, especially for top performing methods.

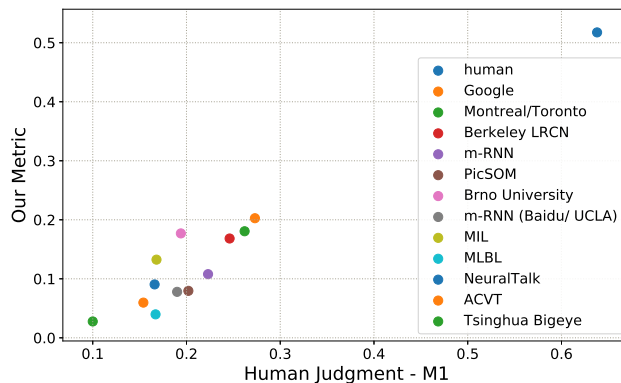


Figure 7. Our metric vs. human judgment on COCO validation set. Our metric is able to reflect most of the rankings from human judgment correctly, especially for top performing methods.

## 5. Conclusion and Future Work

In this paper, we have proposed a novel learning based evaluation metric for image captioning that is trained to act like a human critique to distinguish between human-written and machine-generated captions while also being flexible to adapt to targeted pathological cases. Further, we have shown how to use data sampling and augmentation to successfully train a metric that behaves robustly against captions generated from pathological transformations. From extensive experimental evaluations, we have demonstrated that the proposed metric is robust and correlates better to human judgments than previous metrics. In conclusion, the proposed metric could be an effective complementary to the existing rule-based metrics, especially when the pathological cases are easy to generate but difficult to capture with traditional hand-crafted metrics.

In this study, we have not taken different personalities among human annotators into consideration. Different human personalities could give rise to different types of human captions. One direction of future work could aim to capture the heterogeneous nature of human annotated captions and incorporate such information into captioning evaluation. Another direction for future work could be training a caption generator together with the proposed evaluation metric (discriminator) in a generative adversarial setting. Finally, gameability is definitely a concern, not only for our learning based metric, but also for other rule-based metrics. Learning to be more robust to adversarial examples is also a future direction of learning based evaluation metrics.

**Acknowledgments.** This work was supported in part by a Google Focused Research Award, AWS Cloud Credits for Research and a Facebook equipment donation. We would like to thank the COCO Consortium for agreeing to run our code on entries in the 2015 COCO Captioning Challenge.



## References

- [1] The coco 2015 captioning challenge. <http://mscoco.org/dataset/#captions-challenge2015>. 2, 7
- [2] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al. Tensorflow: A system for large-scale machine learning. In *OSDI*, 2016. 5
- [3] P. Anderson, B. Fernando, M. Johnson, and S. Gould. Spice: Semantic propositional image caption evaluation. In *ECCV*, 2016. 1, 2, 7, 8
- [4] S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio. Generating sentences from a continuous space. In *CoNLL*, 2016. 2
- [5] M. Charikar, K. Chen, and M. Farach-Colton. Finding frequent items in data streams. In *ICALP*, 2002. 3
- [6] T.-H. Chen, Y.-H. Liao, C.-Y. Chuang, W.-T. Hsu, J. Fu, and M. Sun. Show, adapt and tell: Adversarial training of cross-domain image captioner. In *ICCV*, 2017. 2
- [7] X. Chen, H. Fang, T.-Y. Lin, R. Vedantam, S. Gupta, P. Dollár, and C. L. Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv*, 2015. 2
- [8] B. Dai, S. Fidler, R. Urtasun, and D. Lin. Towards diverse and natural image descriptions via a conditional gan. In *ICCV*, 2017. 2
- [9] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 3
- [10] A. Fukui, D. H. Park, D. Yang, A. Rohrbach, T. Darrell, and M. Rohrbach. Multimodal compact bilinear pooling for visual question answering and visual grounding. In *EMNLP*, 2016. 3
- [11] Y. Gao, O. Beijbom, N. Zhang, and T. Darrell. Compact bilinear pooling. In *CVPR*, 2016. 3, 5
- [12] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NIPS*, 2014. 2
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3, 5
- [14] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 1997. 3
- [15] M. Hodosh, P. Young, and J. Hockenmaier. Framing image description as a ranking task: Data, models and evaluation metrics. *J. Artif. Int. Res.*, 47(1):853–899, May 2013. 5, 7, 8
- [16] D. J. Im, C. D. Kim, H. Jiang, and R. Memisevic. Generating images with recurrent adversarial networks. *arXiv*, 2016. 2
- [17] A. Kannan and O. Vinyals. Adversarial evaluation of dialogue models. *arXiv*, 2017. 2
- [18] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015. 4, 5
- [19] D. Kingma and J. Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. 5
- [20] M. D. A. Lavie. Meteor universal: Language specific translation evaluation for any target language. In *ACL*, 2014. 1, 2
- [21] J. Li, W. Monroe, T. Shi, A. Ritter, and D. Jurafsky. Adversarial learning for neural dialogue generation. In *EMNLP*, 2017. 2
- [22] X. Liang, Z. Hu, H. Zhang, C. Gan, and E. P. Xing. Recurrent topic-transition gan for visual paragraph generation. 2017. 2
- [23] C.-Y. Lin. Rouge: A package for automatic evaluation of summaries. In *ACL Workshop*, 2004. 1, 2
- [24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5
- [25] S. Liu, Z. Zhu, N. Ye, S. Guadarrama, and K. Murphy. Improved image captioning via policy gradient optimization of spider. In *ICCV*, 2017. 1, 2
- [26] R. Lowe, M. Noseworthy, I. Serban, N. Angelard-Gontier, Y. Bengio, and J. Pineau. Towards an automatic turing test: Learning to evaluate dialogue responses. In *ACL*, 2017. 2
- [27] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu. Bleu: a method for automatic evaluation of machine translation. In *ACL*, 2002. 1, 2
- [28] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 3, 5
- [29] N. Pham and R. Pagh. Fast and scalable polynomial kernels via explicit feature maps. In *KDD*, 2013. 3
- [30] R. Shetty, M. Rohrbach, L. A. Hendricks, M. Fritz, and B. Schiele. Speaking the same language: Matching machine to human captions by adversarial training. In *ICCV*, 2017. 2
- [31] R. Vedantam, C. Lawrence Zitnick, and D. Parikh. Cider: Consensus-based image description evaluation. In *CVPR*, 2015. 1, 2
- [32] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. In *CVPR*, 2015. 4, 5
- [33] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015. 4, 5