VIFIDEL: Evaluating the Visual Fidelity of Image Descriptions

Pranava Madhyastha[⇔], Josiah Wang[⇔] and Lucia Specia Department of Computing Imperial College London, UK {pranava, josiah.wang, l.specia}@imperial.ac.uk

Abstract

We address the task of evaluating image description generation systems. We propose a novel image-aware metric for this task: VIFIDEL. It estimates the faithfulness of a generated caption with respect to the content of the actual image, based on the semantic similarity between labels of objects depicted in images and words in the description. The metric is also able to take into account the relative importance of objects mentioned in human reference descriptions during evaluation. Even if these human reference descriptions are not available, VIFIDEL can still reliably evaluate system descriptions. The metric achieves high correlation with human judgments on two well-known datasets and is competitive with metrics that depend on and rely exclusively on human references.

1 Introduction

A popular task at the intersection of computer vision and natural language is *image description generation* (IDG), i.e. the task of generating as output a sentence describing the visual content of a given input image. While a variety of methods have been proposed for this task (Kulkarni et al., 2011; Li et al., 2011; Vinyals et al., 2015), its evaluation is still an understudied problem. Evaluation of IDG is currently performed in two ways: (i) human judgment; (ii) automatic metrics. Human judgments evaluate either the overall quality of descriptions or specific criteria in isolation (relevance, fluency, etc.). Such methods, however, can be subjective and expensive to scale.

Automatic metrics address the scalability issue by comparing candidate descriptions against human-authored reference descriptions. These metrics conflate various criteria implicitly into a single evaluation assumption, i.e. a good description is one that is similar to one or more humanauthored descriptions, presuming that these gold descriptions are fluent, correct and relevant to the image. Existing automatic metrics are thus useful for measuring the quality of descriptions as a whole, but this makes it difficult for the specific capabilities of IDG systems to be inspected.

We argue that a fine-grained metric measuring specific criteria would be more useful in understanding how an IDG system is better than another. We focus on one such criterion, visual fi $delity^1$. This criterion aims to measure how faithful a description is with respect to what is depicted in the image (i.e. systems should be rewarded for describing elements depicted in the image and penalised for describing things that are not depicted). For that, we propose to take image content into account when evaluating descriptions, in contrast to previous work (§2) that rely solely on words in the reference descriptions. Given that most datasets are crowd-sourced, reference descriptions may not always accurately describe the image (describing non-depicted objects, not mentioning relevant objects, e.g. Figure 1). For reliable evaluation, multiple references are needed. Image annotations (with objects, attributes, relations, etc.), on the other hand, are arguably more general and less ambiguous for evaluating visual fidelity: they require a single annotation per image, and are not affected by language or style preferences. To our knowledge, no existing metric for IDG has images factored explicitly into the evaluation process.

Our **main contribution** is therefore an automatic evaluation metric for IDG that measures the fidelity of image descriptions with respect to the image, using information derived from images directly as 'reference' (Figure 2). We name

 $[\]Leftrightarrow$ Pranava Madhyastha and Josiah Wang contributed equally as joint first authors to this work.

¹Also referred to as *relevance*, *faithfulness* or *correctness*.



(a) 'What is this??'



(b) 'so what else is she supposed to do?'

Figure 1: Examples of inadequate references taken from (a) PASCAL-50S and (b) MSCOCO.

it VIFIDEL (VIsual Fidelity for Image Description EvaLuation). VIFIDEL can be used (i) based only on images as reference (no textual references) $(\S3.2)$, or (ii) in conjunction with textual references to take into account the relevant image content people describe ($\S3.3$). In addition, VIFIDEL performs matching of images and text in an embeddings space, thus drawing both modalities together semantically while avoiding the pitfalls of mainstream metrics that rely on exact or approximate string matching. This is done by building on the Word Mover's Distance (WMD) metric, which measures the distance between two texts in a word embeddings space. Another contribution is the extension of WMD to allow for multiple references to be used to model object importance, i.e. an approach for consensus within WMD. We evaluate the performance of VIFIDEL against human judgments on two popular IDG datasets ($\S4$).

2 Background

Various IDG metrics have either been adapted from other fields or proposed specifically for IDG. Examples of the former include BLEU (Papineni et al., 2002) and Meteor (Denkowski and Lavie, 2014) from machine translation evaluation, and ROUGE (Lin, 2004) (more specifically ROUGE_L (Lin and Och, 2004)) from text summarisation evaluation. Metrics designed specifically for evaluating image descriptions include CIDEr (Vedantam et al., 2015), SPICE (Anderson et al., 2016) and BAST (Ellebracht et al., 2015).

One main weakness of most metrics (BLEU, ROUGE, Meteor, CIDEr) is that they rely on exact string matching to measure the surface-level, *n*gram overlap between candidate texts and human references. This can result in data spasity problems, especially with limited references. Meteor partially addresses this by matching synonyms from dictionaries and paraphrase tables, but it is constrained to the availability of such dictionaries,



Figure 2: An illustration of VIFIDEL as a visual fidelity metric for IDG. Bright yellow arrows indicate that the word is semantically similar to the object according to Word Mover's Distance, while darker arrows show that they are less related (larger distance).

making it hard to scale to different languages.

Metrics like SPICE and BAST address the issue of exact string matching by measuring similarity on a semantic level. However, these methods rely heavily on linguistic resources such as parsers, semantic role labellers, tailored rules, etc., making evaluation difficult to scale or adapt to different languages and domains. Word Mover's Distance (WMD) (Kusner et al., 2015) has also been proposed as an IDG metric (Kilickaya et al., 2017). WMD finds optimal alignments between word embeddings in candidate and reference descriptions instead of performing *n*-gram matching to address data sparseness issues. VIFIDEL is inspired by WMD, but goes beyond using reference texts by comparing candidate texts against image content ($\S3.2$).

Using images for IDG evaluation. As far as we are aware, no previous work uses images for evaluating IDG. The closest related work is that by Wang and Gaizauskas (2015), who propose an f-measure-based metric to evaluate the task of selecting relevant object instances to be mentioned in a description. The metric computes the overlap between selected object instances and objects mentioned in references, averaged over multiple references. The averaging process implicitly captures consensus over which objects should be mentioned $(\S3.3)$, i.e. objects mentioned in more references should be more important than those mentioned in fewer. Their work, however, requires manual correspondence annotations between bounding box instances and object mentions in descriptions. Our proposed method leverages word embeddings to circumvent the need for exact correspondence annotations.

Complementarity of metrics. Kilickaya et al. (2017) report that the *n*-gram based metrics, the semantic graph based SPICE, and the embedding-based WMD capture complementary information, and that linearly combining Meteor, SPICE and WMD gives a better correlation score against human judgments. Similarly, Liu et al. (2017) optimise image description generation to capture both semantic faithfulness and syntactic fluency by combining metrics with complementary properties. However, the combination weights need to be engineered towards the task.

3 VIFIDEL

In this section we describe the VIFIDEL metric. It is inspired by Word Mover's Distance (WMD) (§3.1), which measures the distance between two documents in a word embedding space. VIFIDEL however explicitly incorporates semantic information derived from images (§3.2), which can be used with or without reference descriptions (§3.3), using word embeddings as a bridge for matching content in images to content in textual descriptions. In contrast to WMD which compares pairs of documents, VIFIDEL also allows for multiple references to be taken into account with a consensus-based approach.

3.1 Word Mover's Distance (WMD)

WMD (Kusner et al., 2015) makes use of word vector relationships between word embeddings to compute the distance between two text documents. WMD captures the minimal distance required to move words from the first document to words in the second document.

Let $X \in \mathcal{R}^{N \times K}$ be a matrix, with Kdimensional word embeddings for a vocabulary of N words. Let $x_i \in \mathcal{R}^K$ be a K-dimensional word embedding vector for word i. A document Δ is represented as an N-dimensional normalised bagof-words (BOW) vector, $d^{\Delta} = (d_1^{\Delta}, d_2^{\Delta}, ..., d_N^{\Delta})$, where d_i^{Δ} is the normalised frequency of word i occurring in document Δ . Stop words are removed from documents; only content words are retained. Kusner et al. (2015) state that stop words are generally less relevant for capturing semantic similarity between documents, especially for bagof-words representations.

As a measure of word-level dissimilarity, Kusner et al. (2015) propose the word travel cost, that is the cost of moving from word i to word j, using the Euclidean distance between the embeddings corresponding to words. More precisely, the cost is defined as:

$$c(i,j) = \|x_i - x_j\|_2^p, \qquad (1)$$

where p is usually 1 or 2 (we set p=2). This allows documents with many closely related words to have smaller distances than documents with very dissimilar words.

To measure distances between two documents α and β , WMD defines a transport matrix $T \in \mathcal{R}^{N \times N}$, where T_{ij} contains information about the proportion of word i in d^{α} that needs to be transported to word j in d^{β} . Formally, WMD computes T that optimises:

$$WMD(d^{\alpha}, d^{\beta}) = \min_{T \ge 0} \sum_{i,j=1}^{N} T_{ij} c(i,j) \quad (2)$$

such that: $\sum_{j=1}^{N} T_{ij} = d_i^{\alpha}$ and $\sum_{i=1}^{N} T_{ij} = d_j^{\beta}$, $\forall i, j$. Here, the normalised bag-of-words distribution of the documents d^{α} and d^{β} contains a combined vocabulary from d^{α} and d^{β} resulting in a square transport matrix T of dimensionality $N \times N$.

Kusner et al. (2015) note that WMD is a special case of Earth Mover's Distance (Rubner et al., 2000), popular in the computer vision community, or Wasserstein's Distance (Datta et al., 2008), popular in the optimal transport community.

3.2 Using objects as image information

In this paper, we explore image information in the form of explicit object detections, both using *gold* or *predicted* object instances for a given image. Previous works (Yin and Ordonez, 2017; Madhyastha et al., 2018; Wang et al., 2018) have found explicit object detections to be informative for image description generation. Thus, we base our intuition on the hypothesis that *a thorough and true description of the image should consist of information about objects and their interactions (frequencies, etc.) in the environment.* VIFIDEL has the capacity to capture these.

While objects represent one important type of semantic image information, VIFIDEL can potentially incorporate other semantic image information including attributes, actions, positions, scenes, relations between objects, and more finegrained information such as colour. For this paper, we consider only the frequency of depicted object category instances as semantic information, and regard further enrichment as future work. As mentioned, a motivation for using image information for evaluating image descriptions is that reference descriptions can be subjective, ambiguous and may or may not specify all (and only) the important elements of the image. In fact, they often focus on a subset of the image content. Using object labels can minimise these issues. In addition, as shown in Figure 1, references are not always actual descriptions and can be incorrect.

Another advantage of a metric based on objectlevel information is the cost of collecting data: either objects are predicted (no labelling involved) or, if gold labels are to be used for more reliable results, object annotations can be gathered in more trustworthy ways using a single annotation per image. With descriptions, it has been shown that multiple descriptions per image are needed for reliable evaluation (Vedantam et al., 2015).

Figure 2 gives an intuitive illustration of our metric. The top row shows a set of detected object instances in the image. The bottom row shows content words in a system description that are semantically very similar to the detected objects – dog, toys, etc. The description does not mention all detected objects (e.g. it misses *table*) and contains the word *beach* that is not in the image. VIFIDEL aims to capture these discrepancies.

More specifically, VIFIDEL is defined as the similarity between the semantic content in image I and a description S. It is specified by the inverse of the minimum cumulative cost required to move semantic labels (e.g. object categories) from image I to words in the description S. This converts WMD from a distance measure to a similarity measure. Formally:

$$VIFIDEL(I,S) = \exp(-WMD(d^{I}, d^{S})) \quad (3)$$

where d^S is the normalised bag of words representation for description S (§3.1), and d^I is a semantic vector representation for image I, specifically a normalised bag of object category labels (§4.1). WMD is defined in Eq. 2.

In its basic form, VIFIDEL can provide information about the compatibility between an image and a description without using reference descriptions. We show the performance of VIFIDEL in the absence of reference descriptions in §4.

3.3 Modelling object importance with reference descriptions

We now expand the basic version of VIFIDEL to use (one or more) human references when available. Human references allow for capturing the *human-likeness* aspect of descriptions, that is, they capture what humans consider important to be described for a picture (Berg et al., 2012).

VIFIDEL can use human references as additional guidance to determine the importance of object content in a given image. We exploit the fact that each reference may only describe a particular subset of the image content, and we assume that important objects are mentioned more frequently across references than less important ones. Our proposal is similar to CIDEr in that we capture consensus information given a set of human references. However, we significantly differ in that (i) we use the references to explicitly model object importance, instead of directly comparing the candidates against the references; (ii) we perform word matching in a semantic space using word embeddings rather than surface forms.

VIFIDEL also differs from previous approaches using WMD for image description evaluation (Kilickaya et al., 2017), where the metric is only computed using the description and the single closest reference. One of the problems in such an approach is the biased choice of reference: in this case, the reference with the smallest WMD distance from the system description. A better reference may be available, e.g. mentioning more image content, which would lead lower system scores in an unbiased evaluation. This has been a common problem in metrics based on a single reference (Fomicheva and Specia, 2016).

Another contribution in this paper is therefore the incorporation of an object importance model into the WMD framework using human references. Our approach rewards candidate descriptions that mention objects depicted in the image (i.e. faithful to image content), and that the objects are also mentioned frequently across all references (i.e. they mention important objects). Intuitively, the WMD cost function (Eq. 1) is replaced with a weighted Euclidean distance. These weights are derived from human descriptions. While the original cost function captures the faithfulness of candidate descriptions to depicted objects, the weights extend the function such that the cost is lower for words that are mentioned frequently across references (either as an exact match or a semantically similar match), and higher for those mentioned less frequently. The weights are applied to both the object labels from images and words for candidate descriptions.

Formally, let $R^{I} = (R_{1}^{I}, R_{2}^{I}, \dots, R_{M}^{I})$ be a set of M human references for image I. The perimage penalty weight, ρ_{k}^{I} , for a word k (an object label in image I or a content word in a candidate description S^{I}) is computed as:

$$\rho_k^I = \frac{1}{M} \sum_{r=1}^M \left(\frac{1 - \max_{t \in \{R_r^I\}} \cos(x_k, x_t)}{2} \right)$$
(4)

where $\{R_r^I\}$ is the set of content words in the *r*th reference for image *I*, and x_t the word embedding for word *t*. The denominator 2 ensures that w_k^I is always in the range [0, 1].

For each image I, we compute the penalty ρ_k^I for each $k \in \{t | d_t^I > 0\} \cup \{t | d_t^{S_I} > 0\}$, i.e. the union of all labels for objects depicted in I and all content words in the candidate description S^I to be evaluated. Thus, ρ_k^I is the effective cosine distance $(\in [0, 1])$ between each word/object label k and its most similar content word in each human reference, averaged over all references for the image. The averaging process implicitly captures the consensus over which objects should be mentioned. ρ_k^I will be small for words/object labels that can are mentioned across most references (using the exact word or a semantically similar word), and large for those that are mentioned only by a few.

The proposed approach of integrating object importance replaces the cost c(i, j) in Eq. 2 with a weighted cost $c'(i, j|R^I)$ to move from word i to word j given references R^I :

$$c'(i,j|R^{I}) = \|\rho_{i}^{I}x_{i} - \rho_{j}^{I}x_{j}\|_{2}^{p}$$
(5)

The updated Eq. 2 is then used in Eq. 3 to compute a VIFIDEL score weighted by object importance. Figure 3 illustrates a concrete example of VIFIDEL's object importance model using human references, showing how the cost $c'(i, j|R^I)$ is calculated.

4 Experiments

We experiment with VIFIDEL in two datasets: the PASCAL-50S Consensus dataset (§4.2) and human ratings on MSCOCO (§4.3). We compare VIFIDEL against commonly used IDG metrics.

4.1 Visual annotation and detectors

We test the performance of the following metric variants, where **VIFIDEL** *gold* and **VIFIDEL** $_{D500}$ and their union are reported in the main experiments (§4.2 and §4.3), while the remaining are used for an ablation study (§4.5):

- VIFIDEL *gold*: This variant uses gold standard, object-level annotations provided by the respective datasets. For the PASCAL-50S dataset, we use the annotations for 20 pre-defined object categories (*person*, *car*, *cow*, etc.) provided by the PASCAL VOC challenge (Everingham et al., 2015). For MSCOCO, we use annotations for 80 object categories provided by MSCOCO (Lin et al., 2014). In both datasets, the reference descriptions are sourced independent of the image annotations; thus there is no direct correspondence between the visual annotations and the descriptions.
- VIFIDEL $_{D80}$: This variant uses the output of an object detector pre-trained on the MSCOCO dataset, for 80 MSCOCO categories. We use the TensorFlow Object Detection API (Huang et al., 2017) for this purpose². We set 0.6 as the confidence threshold for detected objects.
- **VIFIDEL** $_{D500}$: This variant uses the output of an object detector, pre-trained on the Open Images dataset (Krasin et al., 2017) with bounding box annotations for 545 object categories. Again, we use the Tensor-Flow Object Detection API³, and set the confidence threshold to 0.4.
- VIFIDEL gold ∪ D500: We combine the outputs of the gold annotation and D500 detector and use unique object labels from the combination.
- VIFIDEL $_{D80 \cup D500}$: We combine the outputs of the D80 and D500 detectors.

In this paper, we use only the output labels of the detectors, and represent the content of an image I as a vector of normalised frequencies over object labels, d^I . A discussion on how the performance of the metric can vary according to the quality of the objects available is given in S4.5. The ideal setting would count on a comprehensive list of objects given by humans.

²faster_rcnn_inception_resnet_v2_atrous_ coco from https://github.com/tensorflow/ models/blob/master/research/object_ detection/g3doc/detection_model_zoo.md.

³ faster_rcnn_inception_resnet_v2_atrous_ oid.



Figure 3: Illustration of how object importance based on human references is computed and integrated into VIFIDEL. In this example, we compute the weights ρ_k^I for *encyclopedias*, *cat* and *books*. The most similar word to *encyclopedias* in each reference description include *books*, *notebook* and *binder* according to the cosine distance between their word embeddings. These are averaged to obtain a consensus penalty score. The word *cat* has a low penalty score because it has either an exact match or a very close semantic match (*kitten*) to all references. The penalty scores are then used as weights to compute the cost c'(book, encyclopedias|I) between the object label *book* from the image and the word *encyclopedias* in the candidate description.

4.2 Accuracy on PASCAL-50S

In this section, we focus on the PASCAL-50S dataset and tackle the binary forced-choice task of predicting: *"which description is more similar to A: B or C?"*, as proposed by Vedantam et al. (2015). We focus on the variant comparing two machine generated captions. The dataset contains multiple crowdsourced image description for each of 1,000 images from the UIUC PASCAL dataset (Rashtchian et al., 2010).

Evaluation of system outputs as relative rankings has long been established as the best practice in many fields where language outputs are produced and no single correct output exists. The WMT yearly evaluation campaigns for machine translation (Bojar et al., 2016), for example, have argued that relative ranking leads to more reliable judgments than absolute scores. We therefore consider our findings on this dataset as the most important.

For the binary forced-choice task, Vedantam et al. (2015) collected 48 descriptions A per image, and formed pairs of descriptions B and C from machine generated descriptions and/or the remaining two human descriptions. This corresponds to the (**MM**) setting of the dataset: comparing two machine generated descriptions. Arguably, this is most interesting subtask from a practical point of view, which is the setting in which evaluation is generally performed.

The dataset thus consists of 1,000 (B,C) pairs. Crowd-sourced gold standard annotations were

	References					
	0	1	5	48		
BLEU ₁	-	0.58	0.61	0.62		
$BLEU_4$	-	0.56	0.59	0.60		
\mathbf{ROUGE}_L	-	0.58	0.61	0.64		
METEOR	-	0.62	0.66	0.68		
CIDEr	-	0.60	0.66	0.68		
SPICE	-	0.65	0.69	0.69		
WMD _{best}	-	0.66	0.70	0.70		
WMD _{worst}	-	0.66	0.66	0.66		
LM	0.54	0.54	0.54	0.54		
VIFIDEL gold	0.68	0.69	0.70	0.71		
VIFIDEL D500	0.69	0.71	0.72	0.71		
VIFIDEL +LM	0.69	0.70	0.71	0.71		
VIFIDEL +CIDEr	0.69	0.71	0.72	0.72		

Table 1: Accuracy of VIFIDEL on the PASCAL-50S binary forced-choice task using 0, 1, 5 and 48 references, comparing two machine generated descriptions.

provided for each 48×1000 (*A*,*B*,*C*) triplet, and the 'consensus' binary label per (*B*,*C*) pair was obtained by majority vote across 48 references. We also provide results on the other splits in the Appendix A.

Results and discussion Table 1 presents the accuracies for the binary forced-choice task for VIFIDEL, compared to the most commonly used IDG metrics: BLEU₁, BLEU₄, ROUGE_L, ME-TEOR, CIDEr-D, and SPICE⁴. We also report results with the standard WMD as the metric and

⁴We use the official COCO (https://github. com/tylin/coco-caption) and SPICE (https:// github.com/peteanderson80/SPICE) scripts.



Figure 4: Accuracy of consensus predictions on PASCAL-50S, across increasing number of references, for the Machine-Machine (MM) subset.

an RNN language model trained on training captions from MSCOCO (Chen et al., 2015). Note that standard WMD is already a strong metric, but it relies on references, like all other metrics in the table and different from VIFIDEL. VIFIDEL only uses references to up-weight or down-weight the matches between image objects and words in the captions. Notable is the metric's performance based *solely* on image information: VIFIDEL can distinguish between a correct and incorrect description - the difficult task of differentiating between two machine generated descriptions, agreeing with human judgments more often than any other IDG metric that uses a reference. With one reference description, VIFIDEL gave the highest accuracy, suggesting that image-side information is indeed helpful for evaluating the quality of image descriptions when comparing two machine generated descriptions.

Even when we take more references into account, VIFIDEL is as good or better than reference-based metric. With 5 references (a feasible number from a practical perspective), VIFIDEL achieves the highest accuracy with a slight improvement in score over zero or one reference. Thus, the object weighting scheme can be seen to help VIFIDEL focus on which objects are more important. With 48 references (the maximum possible, only available in PASCAL-50S), VIFIDEL is still more accurate than existing metrics, showing that visual fidelity is an important factor for rating machine generated descriptions.

Figure 4 depicts accuracies over an increasing number of reference descriptions, starting from zero references (only defined for VIFIDEL). We note that VIFIDEL is more stable and consistently *outperforms* other metrics for all numbers of references. VIFIDEL $_{D500}$ has a very slight advantage over VIFIDEL $_{gold}$, most likely because the visual information is richer in the former.

Overall, we conclude that measuring visual fidelity is important especially for ranking two machine generated descriptions, arguably the most important evaluation setting, and that VIFIDEL can measure visual fidelity by explicitly using information derived from images, particularly when few or no references are available.

4.3 Correlation with human judges

We measure correlation with human judgments on the MSCOCO portion of the COMPOSITE dataset (Aditya et al., 2018). The dataset contains 2,007 candidate images from the MSCOCO dataset (Lin et al., 2014) with annotations from AMT workers. These annotations consist of judgments on a scale of 1 (low) to 5 (high) regarding (i) relevance of a description to the image and (ii) thoroughness of a description given an image. Candidate descriptions were sampled from human references and image description systems. To make our settings closer to those of real system evaluations, we consider only system generated descriptions as candidates and evaluate the Spearman's correlation between the automated metrics and human judgments. As mentioned in §4.2, absolute human judgments, especially for subjective tasks such as this, can be very subjective and therefore less reliable, especially when a single judgment is collected per description.

Results and discussion We summarise our results in Table 2. For these experiments, we report $gold \cup D500$ detectors for object information, as this was the best overall performance (non-gold and other variants can be seen in §4.5). Our key finding is that VIFIDEL using no reference descriptions obtains comparable (albeit still lower) correlation to metrics like BLEU and ROUGE with only one reference description. The gap between VIFIDEL and such metrics is the inability of the former to capture fluency, since it relies essentially on bag of word embeddings.

4.4 Combining VIFIDEL with other metrics

As mentioned, VIFIDEL is a fine-grained metric that exclusively evaluates the visual fidelity of an image description. As such, by design, an IDG system may achieve high VIFIDEL scores sim-

		5 Refs		1 Ref	0 Ref		
Metric	Relevance	Thoroughness	Relevance	Thoroughness	Relevance	Thoroughness	
BLEU1	0.26	0.25	0.23	0.20	-	-	
\mathbf{BLEU}_4	0.26	0.24	0.22	0.20	-	-	
\mathbf{ROUGE}_L	0.28	0.26	0.24	0.22	-	-	
METEOR	0.30	0.27	0.25	0.22	-	-	
CIDEr	0.32	0.28	0.26	0.24	-	-	
SPICE	0.35	0.31	0.27	0.27	-	-	
WMD _{Best}	0.30	0.26	0.28	0.26	-	-	
WMD _{Worst}	0.29	0.26	0.28	0.25	-	-	
LM	-0.10	-0.16	-0.10	-0.16	-0.10	-0.16	
VIFIDEL no refs.gold UD500	0.22	0.21	0.22	0.21	0.21	0.22	
VIFIDEL $gold \cup D500$	0.30	0.27	0.29	0.27	-	-	
VIFIDEL _{no refs} +LM	0.21	0.21	0.21	0.21	0.21	0.21	
VIFIDEL+CIDEr	0.29	0.29	0.29	0.28	-	-	

Table 2: Spearman's rank correlation coefficient between automatic metrics and human judgment scores over the MSCOCO portion of COMPOSITE dataset.

ply by listing all objects depicted in the image. We do not consider this an issue, since we are in effect evaluating what the image description describes (visual fidelity), rather than how the image is being described (fluency). The latter can instead be evaluated separately with a metric designed specifically for the purpose. Thus, a system that simply lists all objects will result in a high VIFIDEL score but a low fluency score. This is in line with our vision of evaluating the specific capabilities of IDG systems to clearly understand how one system is better than another. While this is not the aim of the paper, we also demonstrate how VIFIDEL can also be combined with fluencybased metrics to evaluate image descriptions as a single conflated metric. More specifically, we explore combining VIFIDEL with two different fluency-based strategies: an RNN language model and CIDEr (other metrics are possible, CIDEr provides a good compromise between performance and efficiency). In both cases, we simply averaged the scores of the two metrics. The results are shown at the bottom of Tables 1 and 2. On average, the addition of the fluency-based metric is complementary. VIFIDEL +CIDEr is better performing than VIFIDEL +LM]. This is expected as LM only provides perplexity scores given a description, while CIDEr explicitly measures quality against references. We note that better combinations can potentially be achieved by learning weights for a weighted average and optimising the training towards fluency.

4.5 Ablation studies

Effect of object detectors and frequency counts: Here we study the effect of using different ob-

	Wit	hout Frequency	With Frequency			
Object Detectors	Relevance Thoroughness		Relevance	Thoroughness		
gold	0.25	0.24	0.25	0.23		
D80	0.24	0.23	0.24	0.22		
D500	0.25	0.24	0.25	0.23		
$gold \cup D500$	0.29	0.27	0.28	0.26		
$D80 \cup D500$	0.28	0.26	0.27	0.26		

Table 3: Ablation experiments on VIFIDEL on the MSCOCO part of the COMPOSITE dataset with different object detectors and with vs. without frequency counts for objects.

ject detectors. We also investigate the contribution of frequency counts in d^{I} , by binarising the frequency counts to indicate only the presence and absence of the objects. The hypothesis is that the number of object instances may be useful for evaluating visual fidelity. The results are summarised in Table 3. The combination of gold and D500object detectors performed the best for the dataset. The gold object information is only slightly better than D80 prediction-based object information. Interestingly, D500, which is more fine-grained, performs as well as the 80-category *gold* object information. Using a binarised d^{I} seemed to give comparable correlation, perhaps even with a marginal edge over its frequency-based counterpart. We postulate that this could be because frequency counts are likely to be mentioned in the descriptions via quantifiers and other morphological and typological variants, which cannot be easily mapped to the frequency of detected objects.

Effect of number of detected objects: We used a pre-trained captioning system (Anderson et al., 2018) to generate captions on a sample of images from the MSCOCO validation set that have different gold object annotations. We present two ex-

VIFIDEL	Objects	Caption	References
0.79	truck	a small truck sitting on top of a field	 A orange tractor sitting on top of a lush green field. A snowplow truck with two snowplows on it Farm equipment truck in a field near a road. an image of a truck with the scaffolding The yellow earth mover sits in the field in front of the pole.
0.70	person, car, backpack, umbrella, handbag, bottle, wine-glass, cup fork, knife, spoon, bowl, broccoli, chair,	a table full of people at the restaurant	 A group of people sitting around a wooden table with food. The nine people smile as they sit at a dinner table. A table full of people that are eating at a restaurant. Co-workers often get together after a long day at work. A group of people that are sitting around a table.

Table 4: Ablation study on the number of detected objects: one detected object and fifteen detected objects.



Figure 5: VIFIDEL score changes with different numbers of object annotations.

amples in Table 4. In the first example the image contains only one object. Here VIFIDEL relies both on the object and on the semantic similarity between the caption and references. In the second example there are fifteen objects, however some are more important than others for describing the images. VIFIDEL gives higher importance to objects that are mentioned in the references. In Figure 5 we further explore the third example from our ablation studies by computing VIFIDEL for different subsets of object annotations, ranging from one object to fifteen object annotations. We see that, as the number of objects increase from one to two, the VIFIDEL score also increases, given that these objects are also mentioned in the system caption. However, with more objects the scores go down, since these are not mentioned in the caption, but the decrease is gradual, even with 15 objects. This happens because such additional objects do not seem too relevant to humans, as these are mostly not mentioned in the references. Therefore, the scores from VIFIDEL change with both number of annotations and the reference descriptions, while metrics like METEOR (0.354)and SPICE (0.320) would remain constant.

Effect of word representations: We also studied the effect of various pre-trained embeddings and found that the pre-trained model of *word2vec* 300-dimensional *CBOW* embeddings (Mikolov et al., 2013) is slightly better than *GLoVe* (Pennington et al., 2014) and *FastText* embeddings (Joulin et al., 2017). This could be because of the amount of data on which these were trained. *FastText* embeddings had similar performance as word2vec embeddings even when only trained on the Wikipedia as corpus. For consistency, we used the *word2vec* embeddings pre-trained on Google News.

5 Conclusions

We have introduced a new metric for image description evaluation that goes beyond comparing descriptions to human references and is explicitly based on object-level image information. Our hypothesis is that the use of image information provides a more reliable pathway for measuring the fidelity of a description for a given image. Further, the metric relies on off-the shelf object detectors and word-embeddings and computes the scores in a semantic space. Our analysis on two of the most widely used datasets for metric comparison shows that our metric correlates well with human judgments, and is particularly well suited when few or no reference description is available. The metric performs comparatively for gold and predicted annotations on objects and is lightweight in terms of dependency on linguistic resources. Our implementation of VIFIDEL can be accessed from: https://github.com/ ImperialNLP/vifidel

Acknowledgements

The authors thank the anonymous reviewers and area chairs for giving up their time to provide useful feedback on an earlier draft of this paper. This work was supported by the MultiMT project (H2020 ERC Starting Grant No. 678017). This work was also supported by the MMVC project, via an Institutional Links grant, ID 352343575, under the Newton-Katip Celebi Fund partnership. The grant is funded by the UK Department of Business, Energy and Industrial Strategy (BEIS) and Scientific and Technological Research Council of Turkey (TÜBİTAK) and delivered by the British Council. For further information, please visit http://www.newtonfund.ac.uk.

References

- Somak Aditya, Yezhou Yang, Chitta Baral, Yiannis Aloimonos, and Cornelia Fermüller. 2018. Image understanding using vision and reasoning through scene description graph. *Computer Vision and Image Understanding*, 173:33–45.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. SPICE: Semantic propositional image caption evaluation. In *Proceedings* of the European Conference on Computer Vision (ECCV), volume 9909 of Lecture Notes in Computer Science, pages 382–398, Amsterdam, The Netherlands. Springer International Publishing.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 6077–6086, Salt Lake City, UT, USA. IEEE.
- Alexander C. Berg, Tamara L. Berg, Hal Daumé III, Jesse Dodge, Amit Goyal, Xufeng Han, Alyssa Mensch, Margaret Mitchell, Aneesh Sood, Karl Stratos, and Kota Yamaguchi. 2012. Understanding and predicting importance in images. In Proceedings of the IEEE Conference on Computer Vision & Pattern Recognition (CVPR), pages 3562–3569, Providence, RI, USA. IEEE.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, Matteo Negri, Aurelie Neveol, Mariana Neves, Martin Popel, Matt Post, Raphael Rubino, Carolina Scarton, Lucia Specia, Marco Turchi, Karin Verspoor, and Marcos Zampieri. 2016. Findings of the 2016 conference on machine translation. In *Proceedings of the First Conference on Machine Translation*, pages 131– 198, Berlin, Germany. Association for Computational Linguistics.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollr, and C. Lawrence Zitnick. 2015. Microsoft COCO captions: data collection and evaluation server. *CoRR*, abs/1504.00325.
- Ritendra Datta, Dhiraj Joshi, Jia Li, and James Z. Wang. 2008. Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys*, 40(2):5:1–5:60.
- 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. Association for Computational Linguistics.
- Lily D. Ellebracht, Arnau Ramisa, Pranava Swaroop Madhyastha, Jose Cordero-Rama, Francesc

Moreno-Noguer, and Ariadna Quattoni. 2015. Semantic tuples for evaluation of image to sentence generation. In *Proceedings of the Fourth Workshop on Vision and Language*, pages 18–28, Lisbon, Portugal. Association for Computational Linguistics.

- M. Everingham, S. M. A. Eslami, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. 2015. The PASCAL Visual Object Classes challenge: A retrospective. *International Journal of Computer Vision*, 111(1):98–136.
- Marina Fomicheva and Lucia Specia. 2016. Reference bias in monolingual machine translation evaluation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 77–82, Berlin, Germany. Association for Computational Linguistics.
- Jonathan Huang, Vivek Rathod, Chen Sun, Menglong Zhu, Anoop Korattikara, Alireza Fathi, Ian Fischer, Zbigniew Wojna, Yang Song, Sergio Guadarrama, and Kevin Murphy. 2017. Speed/accuracy tradeoffs for modern convolutional object detectors. In *Proceedings of the IEEE Conference on Computer Vision & Pattern Recognition (CVPR)*, pages 3296– 3297, Honolulu, HI, USA. IEEE.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers, pages 427–431, Valencia, Spain. Association for Computational Linguistics.
- Mert Kilickaya, Aykut Erdem, Nazli Ikizler-Cinbis, and Erkut Erdem. 2017. Re-evaluating automatic metrics for image captioning. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers, pages 199–209, Valencia, Spain. Association for Computational Linguistics.
- Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Shahab Kamali, Matteo Malloci, Jordi Pont-Tuset, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanesh Narayanan, and Kevin Murphy. 2017. OpenImages: A public dataset for large-scale multi-label and multiclass image classification. Dataset available from https://storage.googleapis.com/ openimages/web/index.html.
- Girish Kulkarni, Visruth Premraj, Sagnik Dhar, Siming Li, Yejin Choi, Alexander C. Berg, and Tamara L. Berg. 2011. Baby talk: Understanding and generating simple image descriptions. In *Proceedings of the IEEE Conference on Computer Vision & Pattern Recognition (CVPR)*, pages 1601–1608, Colorado Springs, CO, USA. IEEE.

- Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 957–966, Lille, France. PMLR.
- Siming Li, Girish Kulkarni, Tamara L. Berg, Alexander C. Berg, and Yejin Choi. 2011. Composing simple image descriptions using web-scale n-grams. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 220– 228, Portland, OR, USA. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In ACL Workshop on Text Summarization Branches Out. Association for Computational Linguistics.
- Chin-Yew Lin and Franz Josef Och. 2004. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics* (ACL-04), pages 605–612, Barcelona, Spain. Association for Computational Linguistics.
- 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 Proceedings of the European Conference on Computer Vision (ECCV), pages 740–755, Zurich, Switzerland. Springer International Publishing.
- Siqi Liu, Zhenhai Zhu, Ning Ye, Sergio Guadarrama, and Kevin Murphy. 2017. Improved image captioning via policy gradient optimization of SPIDEr. In Proceedings of the IEEE International Conference on Computer Vision (ICCV), pages 873–881, Venice, Italy. IEEE.
- Pranava Madhyastha, Josiah Wang, and Lucia Specia. 2018. End-to-end image captioning exploits multimodal distributional similarity. In Proceedings of the 29th British Machine Vision Conference (BMVC), page 306, Newcastle, UK. BMVA Press.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. In *Proceedings of the 1st International Conference on Learning Representations, ICLR 2013, Workshop Track Proceedings*, Scottsdale, AZ, USA.
- 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, PA, USA. Association for Computational Linguistics.

- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Cyrus Rashtchian, Peter Young, Micah Hodosh, and Julia Hockenmaier. 2010. Collecting image annotations using Amazon's Mechanical Turk. In Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk, pages 139–147, Los Angeles, CA, USA. Association for Computational Linguistics.
- Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. 2000. The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2):99–121.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. CIDEr: Consensus-based image description evaluation. In Proceedings of the IEEE Conference on Computer Vision & Pattern Recognition (CVPR), pages 4566–4575, Boston, MA, USA. IEEE.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE Conference on Computer Vision & Pattern Recognition (CVPR)*, pages 3156–3164, Boston, MA, USA. IEEE.
- Josiah Wang and Robert Gaizauskas. 2015. Generating image descriptions with gold standard visual inputs: Motivation, evaluation and baselines. In Proceedings of the 15th European Workshop on Natural Language Generation (ENLG), pages 117–126, Brighton, UK. Association for Computational Linguistics.
- Josiah Wang, Pranava Madhyastha, and Lucia Specia. 2018. Object counts! Bringing explicit detections back into image captioning. In Proceedings of the North American Chapter of the Association of Computational Linguistics: Human Language Technologies (NAACL HLT), Volume 1 (Long Papers), pages 2180–2193, New Orleans, LA, USA. Association for Computational Linguistics.
- Xuwang Yin and Vicente Ordonez. 2017. Obj2Text: Generating visually descriptive language from object layouts. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing* (*EMNLP*), pages 177–187, Copenhagen, Denmark. Association for Computational Linguistics.

A Accuracy on PASCAL-50S (all groups)

The original PASCAL-50S dataset consisted of four groups: (i) **HC** (Human-Human correct: both correct human descriptions); (ii) **HI** (Human-Human incorrect: one correct human description

and one human description from another random image); (iii) **HM** (Human-Machine: one correct human description and one machine generated description); (iv) **MM** (Machine-Machine: both machine generated descriptions). We summarise the accuracy of VIFIDEL and various evaluation metrics on the PASCAL-50S binary forced-choice task in Tables 5 and 6.

Metric	No references					1 reference
	HC	HI	HM	MM	all	HC HI HM MM all
BLEU ₁	-	-	-	-	-	0.61 0.86 0.88 0.58 0.73
$BLEU_4$	-	-	-	-	-	0.59 0.83 0.80 0.56 0.70
$ROUGE_L$	-	-	-	-	-	0.61 0.87 0.86 0.58 0.73
METEOR	-	-	-	-	-	0.61 0.91 0.89 0.62 0.76
CIDEr	-	-	-	-	-	0.62 0.91 0.85 0.60 0.75
SPICE	-	-	-	-	-	0.58 0.87 0.77 0.65 0.72
VIFIDEL gold	0.57	0.91	0.53	0.68	0.67	0.62 0.96 0.65 0.69 0.73
VIFIDEL D500	0.57	0.93	0.64	0.69	0.71	0.59 0.95 0.74 0.71 0.75

Table 5: Accuracy of VIFIDEL and other IDG metrics on the PASCAL-50S binary forced-choice task using no references and one reference.

Metric	5 references						48 references				
Meule	HC HI HM MM all	HC	HI	HM	MM	all					
BLEU ₁	0.64	0.95	0.91	0.61	0.78	0.63	0.98	0.94	0.62	0.79	
$BLEU_4$	0.62	0.93	0.85	0.59	0.75	0.64	0.97	0.92	0.60	0.78	
$ROUGE_L$	0.66	0.95	0.93	0.61	0.78	0.67	0.98	0.95	0.64	0.81	
METEOR	0.66	0.98	0.94	0.66	0.81	0.65	0.99	0.96	0.68	0.82	
CIDEr	0.67	0.98	0.90	0.66	0.80	0.69	0.99	0.92	0.68	0.82	
SPICE	0.67	0.97	0.89	0.69	0.80	0.62	0.99	0.94	0.69	0.81	
VIFIDEL aold	0.65	0.98	0.66	0.70	0.75	0.64	0.98	0.67	0.71	0.75	
VIFIDEL D500	0.64	0.97	0.75	0.72	0.77	0.63	0.97	0.76	0.71	0.77	

Table 6: Accuracy of VIFIDEL and other IDG metrics on the PASCAL-50S binary forced-choice task using 5 and 48 references.