

# InfoMetIC: An Informative Metric for Reference-free Image Caption Evaluation

Anwen Hu<sup>1</sup>, Shizhe Chen<sup>2</sup>, LiangZhang<sup>1</sup>, Qin Jin<sup>1\*</sup>

<sup>1</sup>School of Information, Renmin University of China

<sup>2</sup>INRIA

{anwenhu, zhangliang00, qjin}@ruc.edu.cn

shizhe.chen@inria.fr

## Abstract


Automatic image captioning evaluation is critical for benchmarking and promoting advances in image captioning research. Existing metrics only provide a single score to measure caption qualities, which are less explainable and informative. Instead, we humans can easily identify the problems of captions in details, e.g., which words are inaccurate and which salient objects are not described, and then rate the caption quality. To support such informative feedback, we propose an **Informative Metric for Reference-free Image Caption evaluation (InfoMetIC)**. Given an image and a caption, InfoMetIC is able to report incorrect words and unmentioned image regions at fine-grained level, and also provide a text precision score, a vision recall score and an overall quality score at coarse-grained level. The coarse-grained score of InfoMetIC achieves significantly better correlation with human judgements than existing metrics on multiple benchmarks. We also construct a token-level evaluation dataset and demonstrate the effectiveness of InfoMetIC in fine-grained evaluation. Our code and datasets are publicly available at <https://github.com/HAWLYQ/InfoMetIC>.

## 1 Introduction

Image captioning aims to automatically generate natural language sentences to describe image contents. Recently, there are significant breakthroughs in image captioning such as attention-based model architectures (Anderson et al., 2018; Pan et al., 2020; Hu et al., 2020, 2021) and vision-and-language pretraining (VLP) (Zhou et al., 2020; Xia et al., 2021; Li et al., 2022b; Xu et al., 2021; Li et al., 2022a). However, as groundtruth image descriptions are extremely diverse and subjective, evaluating the image captioning performance remains a considerable challenge.

The most widely used image captioning metrics such as METEOR (Banerjee and Lavie, 2005),

\*Corresponding Author.



### References:

1. A brown and white dog is running.
2. A dog runs on the snow with a package.
3. There is a hat buried in the snow.

**Candidate 1(C1):** A dog on the snow.

**Candidate 1(C2):** A dog with a hat is on the snow.

**Candidate 2(C3):** A dog with a bag is on the snow.

	METEOR (w/ ref)	SPICE (w/ ref)	CLIP-S (w/o ref)	Ours: InfoMetIC (w/o ref)		
				Text Precision	Vision Recall	Overall
C1	0.27	0.31	0.78	dog✓ 5.71 snow✓	0.62 ✓✓✓	6.33
C2	0.37	0.40	0.76	dog✓ 4.78 hat✗ snow✓	1.27 ✓✓✓	6.05
C3	0.37	0.27	0.84	dog✓ 6.45 bag✓ snow✓	1.69 ✓✓✓	8.14

Figure 1: Comparison of existing metrics and our informative metric (InfoMetIC). ‘w/ ref’ and ‘w/o ref’ mean using references or not.

CIDEr (Vedantam et al., 2015a) and SPICE (Anderson et al., 2016) utilize human-written descriptions of images as references and measure similarities between generated captions and references for evaluation. Such reference-based approaches suffer from two major limitations. Firstly, these metrics mainly evaluate caption quality by n-gram overlaps which fail to measure genuine semantic similarities. Secondly, references require time-consuming annotations and thus there are only a few annotated captions (typically 5) for each image. The limited number of references cannot fully capture image contents, resulting in incorrect penalties when generated captions describe correct novel things that are not mentioned in the references.

To alleviate the above limitations, recent works are more focusing on reference-free metrics, which directly use images instead of reference captions in evaluation. Benefited from the success of VLP on large-scale web data, UMIC (Lee et al., 2021) and CLIP-S (Hessel et al., 2021) leverage VLP models UNITER (Chen et al., 2020) and CLIP (Radford et al., 2021) respectively to calculate relevance scores between generated captions and images. Although they have achieved promising correlations with human judgments, they can only produce an overall score as quality measurement. We humans

instead tend to evaluate captions considering two aspects: 1) whether the caption correctly describes the image content (named *text precision*); and 2) whether the image content is comprehensively described in the caption (named *vision recall*). For example, as shown Figure 1, we can easily tell the “hat” in the second candidate is incorrect, and some salient contents such as “the bag” are not mentioned, and thus form our final evaluation to the caption.

For the purpose of providing explainable and detailed feedbacks, we propose a **Informative Metric** for Reference-free Image Caption evaluation (InfoMetIC). It is built on top of pretrained VLP models to measure fine-grained cross-modal similarities. InfoMetIC is able to point out incorrect semantic words in the caption and unmentioned regions in the image. Based on fine-grained evaluation, it derives text precision and vision recall scores to measure captioning accuracy and completeness respectively. We take a summation of the two scores to rate overall quality of the caption.

Our contributions in this work are three-fold:

- We propose a reference-free informative image captioning metric InfoMetIC. It can provide both coarse-grained scores and detailed token-level scores.
- We automatically construct training examples based on annotations in image caption datasets and design coarse- and fine-grained tasks to train the evaluation model.
- InfoMetIC achieves better correlation with human judgements on multiple benchmarks, as well as on our newly constructed fine-grained caption evaluation benchmark CapTokenEval.

## 2 Related Work

**Reference-only caption evaluation.** This type of evaluation only employs human-written captions as references and measures text similarity as the evaluation score. Most widely used metrics such as BLEU-4 (Papineni et al., 2002), ROUGE-L (Lin, 2004), METEOR (Banerjee and Lavie, 2005), CIDEr (Vedantam et al., 2015a) and SPICE (Anderson et al., 2016) all fall into this category. BLEU-4 calculates the precision of n-gram matches; ROUGE-L measures the recall of the longest common subsequence; METEOR utilizes wordnet-based synonym matching to relieve the shortage of exact word matching; CIDEr introduces tf-idf to re-weight the importance of different n-grams; SPICE converts

captions into scene graphs for similarity comparison. One major limitation of the above metrics is that they cannot properly count synonym matches. To overcome this deficiency, BERT-S (Zhang et al., 2020) leverages learned embeddings from a pre-trained language model BERT (Devlin et al., 2019) to better measure semantic similarities. BERT-S++ (Yi et al., 2020) further improves BERT-S by taking into account the variance of multiple references.

**Reference+image caption evaluation.** As an image is worth a thousands of words, a limited number of references cannot fully cover image contents, making the reference-only caption evaluation less reliable. Therefore, some works combine both references and images to evaluate generated captions. REO (Jiang et al., 2019a) uses a pretrained image-text retrieval model SCAN (Lee et al., 2018) to extract image contextualized caption features for computing relevance, extraneous and omission scores. TIGER (Jiang et al., 2019b) calculates grounding vectors for captions via SCAN to measure similarity, which represent how much captions are grounded in an image. ViLBERTScore (Lee et al., 2020) is similar to BERT-S except that it generates visually-grounded features for each caption token by ViLBERT (Lu et al., 2019). FAIEr (Wang et al., 2021) fuses scene graphs of the image and references as a union scene graph and compares it with the scene graph of generated captions.

**Reference-free caption evaluation.** To alleviate the annotation burden of obtaining references, a few works propose to evaluate image captions without references. UMIC (Lee et al., 2021) fine-tunes a pretrained multimodal transformer UNITER (Chen et al., 2020) by contrastive learning to compute an image-text matching score. CLIP-S (Hessel et al., 2021) directly utilizes image-text similarity from CLIP (Radford et al., 2021) - an image-text matching model trained on large-scale open-domain data. CLIP-S has achieved state-of-the-art evaluation performance. However, these methods only provide single scores which are less informative to evaluate image captions. In this work, we aim to provide more fine-grained feedbacks, not only indicating the captioning quality from precision and recall aspects, but also pointing out detailed mistakes such as incorrect words and unmentioned regions.

## 3 Method

We first introduce our model architecture in Sec 3.1 and then describe the training and inference ap-

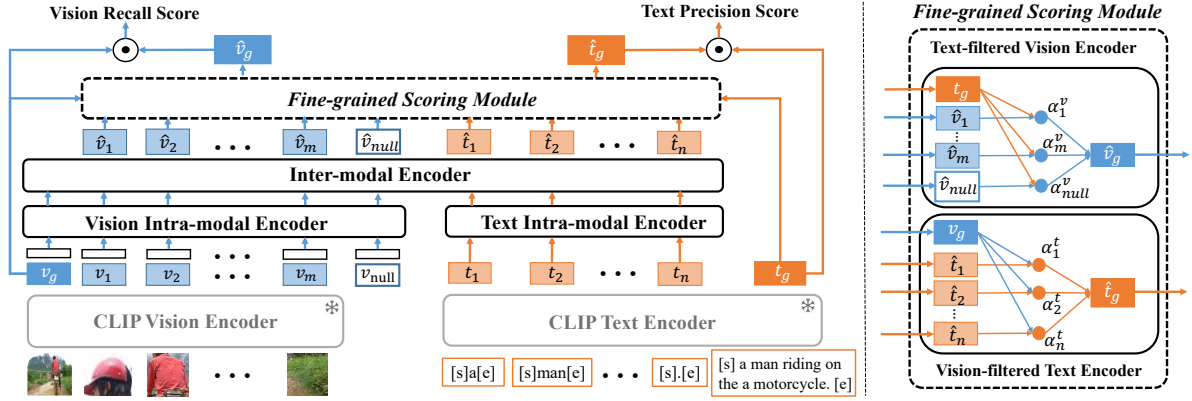


Figure 2: **Left:** the overall architecture of **Informative Metric** for Reference-free Image Caption evaluation (InfoMetIC). **Right:** the detailed structure of the *Fine-grained Scoring Module*.

proaches in Sec 3.2 and Sec 3.3 respectively.

### 3.1 Model Architecture

Figure 2 illustrates the overall framework of our informative evaluation model, which consists of three modules: *Token-level Encoding*, *Intra&Inter Modality Fusion* and *Fine-grained Scoring*. Given an image  $I$  and a caption  $C$  as inputs, the Token-level Encoding module firstly generates a sequence of token-level features to represent the image and caption respectively. Then the Intra&Inter Modality Fusion module captures the intra- and inter-modality relationships. Finally, the Fine-grained Scoring module produces token-level scores for each visual and textual token and derives vision recall, text precision, and overall scores based on the token-level scores.

#### 3.1.1 Token-level Encoding

VLP models have shown superior performance and generalization ability in many vision-and-language tasks (Chen et al., 2020). Therefore, we utilize a state-of-the-art VLP model CLIP to extract token-level image and caption features. To be noted, our method can be adapted to different VLP models.

**Image Token Features.** In order to obtain semantically meaningful image tokens, we use a pretrained object detector to detect region bounding boxes in image  $I$ . We encode each cropped region via CLIP vision encoder to get fine-grained token-level features  $(v_1, \dots, v_m)$ , where  $m$  is the number of detected regions. The whole image is encoded as a global vision feature  $v_g$ . We further utilize a zero vector to represent a vision null token  $v_{null}$ , which aims to align with any texts irrelevant to the image.

**Caption Token Features.** For a caption  $C$ , CLIP

text encoder can generate a global feature  $t_g$  to capture overall semantics of the whole sentence. Although it could also generate a sequence of text token features, these features can overuse the sentence context, which harms fine-grained evaluation. An illustration about the context overuse can be found in Appendix A. Therefore, we encode each token in  $C$  separately as shown in Figure 2 to obtain independent token-level features  $(t_1, \dots, t_n)$ , where  $n$  is the number of text tokens.

#### 3.1.2 Intra&Inter Modality Fusion

In order to learn intra-modal relationships, we utilize two multi-layer transformers (Vaswani et al., 2017) to encode image and text tokens separately. As spatial information is essential to infer relationships across image regions, we apply a linear layer to convert normalized bounding boxes as position features and add them to the initial image token features before fed into the intra-modal transformer. Likewise, we add learnable position features for the text tokens. For visual intra-modal encoding, we concatenate  $v_g$  with  $(v_1, \dots, v_m, v_{null})$  to alleviate possible vision context loss in fine-grained image tokens due to imperfect detection. For textual intra-modal encoding, we directly utilize  $(t_1, \dots, t_n)$  tokens as inputs.

We concatenate the image and text token-level features after intra-modal encoding and utilize an inter-modal encoder to learn correlation between vision and text modalities. The inter-modal encoder is implemented as a multi-layer cross-modal transformer (Chen et al., 2020). We denote the output features for image tokens as  $\hat{V} = (\hat{v}_1, \dots, \hat{v}_m, \hat{v}_{null})$ , output features for text tokens as  $\hat{T} = (\hat{t}_1, \dots, \hat{t}_n)$ .

### 3.1.3 Fine-grained Scoring

The Fine-grained Scoring module aims to predict which text tokens are incorrect and which image tokens are not mentioned. It consists of two cross-modal attention layers, namely Text-filtered Vision Encoder and Vision-filtered Text Encoder as shown in the right of Figure 2. To identify which image tokens are mentioned, we use global text feature  $t_g$  as query and token-level vision features  $\hat{V}$  as key in the cross-modality attention layer to calculate visual token-level scores  $\alpha^v$ :

$$s_i^v = (t_g W_q^v)^T \hat{v}_i W_k^v, \quad (1)$$

$$\alpha^v = \text{Softmax}([s_1^v, \dots, s_m^v, s_{null}^v]). \quad (2)$$

Similarly, to identify which text tokens are incorrect, we use global vision feature  $v_g$  as query and token-level text features  $\hat{T}$  as key to calculate textual token-level scores  $\alpha^t$  by another cross-modality attention layer.

Based on token-level scores, we derive vision recall score and text precision scores to measure the comprehensiveness and accuracy of generated captions respectively. We take visual token-level scores  $\alpha^v$  and token-level vision features  $\hat{V}$  to obtain a text-conditioned vision feature  $\hat{v}_g$  by weighed average as follows:

$$\hat{v}_g = \sum_{k \in \{1, \dots, m, null\}} \alpha_k^v \hat{v}_k. \quad (3)$$

The more image regions are mentioned in a caption, the closer its text-conditioned vision feature should be to the global vision feature  $v_g$ . Thus, we compute the vision recall score as the cosine similarity between  $\hat{v}_g$  and  $v_g$ , represented as  $f^R(I, C) = \cos(\hat{v}_g, v_g) / \tau$ , where  $\tau$  is a learnable temperature parameter. Taking the untrained global vision feature  $v_g$  as the comparison object, our vision recall score implicitly considers the salience of visual information, as illustrated in Appendix B. In a similar way, we can obtain a vision-conditioned text feature  $\hat{t}_g$  and compute a text precision score  $f^P(I, C) = \cos(\hat{t}_g, t_g) / \tau$ . Our overall score is the summation of precision score and recall score:

$$f^O(I, C) = f^R(I, C) + f^P(I, C). \quad (4)$$

## 3.2 Multi-task Learning

To learn fine-grained token-level predictions as well as coarse-grained text precision and vision recall scores, we propose multiple training tasks to jointly optimize our evaluation model.

### 3.2.1 Coarse-grained Score Learning

Given an aligned image-caption pair  $(I, C)$ , we construct negative samples by pairing  $I$  with other captions in the training batch or pairing  $C$  with other images in the batch. Then, we calculate Noisy Contrastive Learning (NCE) loss  $l_r$  based on vision recall scores and  $l_p$  based on text precision scores. The NCE loss  $l_r$  is calculated as follows:

$$l_r = (l_r^i + l_r^c) / 2, \quad (5)$$

$$l_r^i = -\mathbb{E}_{(I, C) \sim B} \log \frac{e^{f^R(I, C)}}{\sum_{C' \in \mathcal{N}_I \cup \{C\}} e^{f^R(I, C')}}, \quad (6)$$

$$l_r^c = -\mathbb{E}_{(I, C) \sim B} \log \frac{e^{f^R(I, C)}}{\sum_{I' \in \mathcal{N}_C \cup \{I\}} e^{f^R(I', C)}}, \quad (7)$$

where  $\mathcal{N}_I$  means a set of negative captions for image  $I$  within the batch  $B$ ,  $\mathcal{N}_C$  means negative images for caption  $C$ . The NCE loss  $l_p$  is similar to Eq (5) but utilizes  $f^P(I, C)$  scores in computation.

**Hard Textual Negatives.** In the above coarse-grained score learning, negative captions for an image are randomly selected from the dataset and usually contains many irrelevant contents with the image. These textual negatives are not hard enough to learn a good vision recall score. Because the model could compute a high recall score for positive pairs by putting high weight to only one rather than all mentioned regions. To address this problem, we further design Hard Textual Negatives (HTN) during coarse-grained score learning. For multiple annotated captions of an image, we consider the one with more semantic words (nouns, verbs, adjectives and adverbs) should get higher vision recall score than the others. Therefore, we treat the other ones as hard textual negatives. The HTN loss  $l_r^h$  is calculated as follows:

$$l_r^h = -\mathbb{E}_{(I, C) \sim B} \log \frac{e^{f^R(I, C)}}{e^{f^R(I, C)} + e^{f^R(I, C^h)}}, \quad (8)$$

where  $C^h$  is a hard textual negative for caption  $C$ .

### 3.2.2 Fine-grained Score Learning

To improve fine-grained evaluation, we design a sequence labeling task called Fine-grained Score learning. We automatically generate supervision signals to learn token-level predictions. For the text part, we prepare labels in a self-supervised manner. Given an image  $I$  and its groundtruth caption  $C$ , we generate a polluted caption  $C'$  by randomly replacing a semantic word with a frequent word

of the same part-of-speech tag. The text sequence label  $Y^t$  for  $(I, C')$  is constructed by setting the polluted word as 0 (incorrect) and other semantic words as 1 (correct). Non-semantic words such as adpositions, conjunctions are excluded in training. For the image part, we make use of existing phrase grounding annotations which align each phrase in a caption with its corresponding bounding boxes in the image. The vision sequence label  $Y^v$  for  $(I, C)$  is constructed by setting all regions mentioned by the caption as 1 and otherwise 0.

We use cross-entropy losses for both textual and visual fine-grained score learning tasks:

$$l_t^{token} = -\frac{1}{n^s} \sum Y^t \log(\alpha^t), \quad (9)$$

$$l_v^{token} = -\frac{1}{m} \sum Y^v \log(\alpha^v), \quad (10)$$

where  $l_t^{token}$  and  $l_v^{token}$  refer to the text-part and vision-part loss respectively,  $\alpha^t$  and  $\alpha^v$  are textual token-level scores and visual token-level scores in Eq (2),  $n^s$  is the number of semantic words.

### 3.3 Inference

Given input pair  $(I, C)$ , we first compute token-level scores  $\alpha^v$  and  $\alpha^t$  for fine-grained prediction with a threshold  $\beta$ . Considering that a caption hardly contains more than 10 semantic words, we set  $\beta$  as 0.1. For the text part, semantic tokens with a score greater than  $\beta$  are judged as correct ones. For the image part, regions with a score greater than  $\beta$  are identified as mentioned ones.

Then we calculate the vision recall, text precision, and overall scores as in Eq (4). We denote our vision recall score  $f^R(I, C)$  as InfoMetIC<sup>R</sup>, text precision score  $f^P(I, C)$  as InfoMetIC<sup>P</sup>, and overall score  $f^O(I, C)$  as InfoMetIC. Furthermore, we combine our overall score with the CLIP similarity:

$$\text{InfoMetIC}^+ = \text{InfoMetIC} + \frac{\cos(v_g, t_g)}{\tau^{clip}} \quad (11)$$

where  $\tau^{clip}$  is the temperature of CLIP.

## 4 Experiment

### 4.1 Experimental Setting

**Training Datasets.** With the training splits of Flickr30k (Young et al., 2014) and MSCOCO (Lin et al., 2014) datasets, we construct 715,662 image-caption pairs for general coarse-grained score learning, and 611,105 triplets with hard textual negatives. For fine-grained score learning, we construct

512,000 samples from MSOCO and Flickr30k for the text part training and 178,689 samples from Flickr30k for the vision part training.

**Implementation Details.** We use CLIP(ViT-B/32) for token-level encoding. The image regions are detected by the bottom-up model (Anderson et al., 2018). To remove redundant bounding boxes, we use k-means algorithm to generate 20 clusters among 100 detected regions and select one region per cluster. The details can be found in Appendix C. The maximum length for textual tokens is set as 32. In the intra&inter modality fusion, intra- and inter-modal encoders contain 4 and 2 transformer layers respectively. During training, the batch size is set as 32 and the initial learning rate is set as 1e-4. We iteratively train our model on multiple tasks for 32,000 iterations. The training ratio of coarse- and fine-grained tasks is 3:1. The training takes 5 hours on 4 V100 GPUs.

## 4.2 Coarse-grained Score Evaluation

### 4.2.1 Evaluation Datasets

**Flickr8k-Expert** (Hodosh et al., 2013a) contains 5,644 pairs of images and machine-generated captions. Each pair is scored from 1 (irrelevant) to 4 (well related) by 3 expert annotators.

**Flickr8k-CF** (Hodosh et al., 2013a) consists of 47,830 image-captions pairs. Each pair is judged "yes" or "no" by at least 3 annotators, where "yes" is for good captions. The final score of each pair is determined by the proportion of "yes".

**Composite** (Aditya et al., 2018) contains 3,995 images from MSCOCO, Flickr30K and Flickr8k (Hodosh et al., 2013b). For each image, there are two machine-generated captions and one human-written caption. Every image-caption pair is scored from 1 (irrelevant) to 5 (perfectly related).

**Pascal-50S** (Vedantam et al., 2015b) contains 4,000 triplets, each of which contains an image and two captions. Annotators are asked to judge which caption is better. According to caption types, Pascal-50S is evenly split into 4 subsets: 'HC' means two correct human-written captions; 'HI' means two human-written captions but one is wrong; 'HM' means one human-written caption and one machine-generated caption; 'MM' means two machine-generated captions.

**THUMB 1.0** (Kasai et al., 2022) contains 500 images from MSCOCO. Each image is paired with one human-written caption and four machine-generated captions. For each image-caption pair,

Table 1: Overall score comparison on Flickr8k-Expert (F-Ex), Flickr8k-CF (F-CF), Composite (Com) and Pascal-50S. ‘w/ ref’ means using 4-5 ground-truth references. ‘w/o ref’ means using no reference.

Type	Metric	F-Ex( $\tau_c$ )	F-CF( $\tau_b$ )	Com( $\tau_c$ )	Pascal-50S (accuracy)				
					HC	HI	HM	MM	Mean
w/ ref	BLEU-4	30.8	16.9	30.6	52.5	90.4	63.0	42.3	55.8
	ROUGE-L	32.3	19.9	32.4	55.0	95.3	93.1	58.7	75.5
	METEOR	41.8	22.2	38.9	59.0	97.7	93.9	62.0	78.2
	CIDEr	43.9	24.6	37.7	53.7	98.1	90.8	63.7	76.6
	SPICE	44.9	24.4	40.3	56.9	96.3	87.1	66.4	76.7
	BERT-S	39.2	22.8	30.1	54.4	96.1	94.3	56.4	75.3
	BERT-S++	46.7	-	44.9	65.4	98.1	96.4	60.3	80.1
	TIGEr	49.3	-	45.4	56.0	99.8	92.8	74.2	80.7
	ViLBERTScore-F	50.1	-	52.4	49.9	99.6	93.1	75.8	79.6
	FAIEr-4	52.6	35.4	57.7	59.7	<b>99.9</b>	92.7	73.4	81.4
RefCLIP-S	53.0	36.4	55.4	57.9	99.5	96.1	80.8	83.6	
w/o ref	UMIC	46.8	-	56.1	66.1	99.8	98.1	76.2	85.1
	FAIEr-r	50.1	32.4	50.5	-	-	-	-	-
	CLIP-S	51.5	34.4	53.8	60.4	99.4	<b>97.8</b>	77.1	83.7
	CLIP- $S_{tune}$	54.3	36.6	57.3	61.0	99.5	95.9	<b>82.0</b>	84.6
	InfoCLIP	32.6	23.5	15.3	37.3	87.3	58.9	72.9	64.1
	InfoCLIP $_{tune}$	37.7	27.7	24.6	37.3	92.5	62.7	74.7	66.8
	InfoMetIC	54.2	36.3	59.2	69.0	99.8	94.0	78.3	85.3
	InfoMetIC <sup>+</sup>	<b>55.5</b>	<b>36.6</b>	<b>59.3</b>	<b>69.9</b>	99.7	96.8	79.6	<b>86.5</b>

there are a precision score measuring the accuracy of the caption, a recall score assessing how much of the salient information is covered, and a total score measuring the overall quality.

#### 4.2.2 Evaluation Metrics

We follow previous works (Hessel et al., 2021; Vedantam et al., 2015b; Kasai et al., 2022) to evaluate captioning metrics. We use kendall-c correlation ( $\tau_c$ ) on Flickr8k-Expert, kendall-b correlation ( $\tau_b$ ) on Flickr8k-CF, kendall-c correlation ( $\tau_c$ ) on Composite, classification accuracy on Pascal-50s and Pearson correlation ( $\rho$ ) on THUMB 1.0.

#### 4.2.3 Comparison with State of the Arts

We compare InfoMetIC with SOTA methods as well as three strong baselines: CLIP- $S_{tune}$ , InfoCLIP and InfoCLIP $_{tune}$ . CLIP- $S_{tune}$  calculates an overall score as CLIP-S (Hessel et al., 2021) but is fine-tuned on MSCOCO and Flickr30k. InfoCLIP directly uses CLIP to perform fine-grained scoring like InfoMetIC but removes the Intra&Inter Modality Fusion and parameters in Fine-grained Scoring. InfoCLIP $_{tune}$  is a fine-tuned version of InfoCLIP. More details can be found in the Appendix D.

Table 1 shows the overall score comparison on Flickr8k-Expert, Flickr8k-CF, Composite and Pascal-50S. Our reference-free metric InfoMetIC achieves state-of-the-art correlation with human judgements on Composite and Pascal-50S. It is on par with the strong baseline CLIP- $S_{tune}$

on Flickr8k-Expert and Flickr8k-CF. To be noted, InfoMetIC performs much better than InfoCLIP, which proves the necessity of our model architecture upon CLIP backbones. After combined with CLIP similarity, InfoMetIC<sup>+</sup> further improves performances on all benchmarks.

To separately evaluate the performance of our vision recall score InfoMetIC<sup>R</sup> and text precision score InfoMetIC<sup>P</sup>, we further conduct experiments on THUMB 1.0 in Table 3. **First**, by comparing InfoMetIC<sup>P</sup> and InfoMetIC<sup>R</sup>, InfoMetIC<sup>R</sup> achieves better correlation with human-labeled recall score and InfoMetIC<sup>P</sup> achieves better correlation with human-labeled precision score. This indicates that our InfoMetIC<sup>R</sup> and InfoMetIC<sup>P</sup> indeed evaluates the recall of image contents and the precision of caption respectively. Besides, both InfoMetIC<sup>P</sup> and InfoMetIC<sup>R</sup> surpass the state-of-the-art reference-free metric CLIP-S on total score correlation. **Second**, our overall score InfoMetIC achieves significant boost on total score, which demonstrates that precision and recall are complementary in human’s final evaluation for captions. InfoMetIC<sup>+</sup> slightly improves the total score performance. **Third**, compared with the state-of-the-art reference-based metric RefCLIP-S (Hessel et al., 2021), our InfoMetIC<sup>+</sup> achieves much better recall correlation but lower precision correlation with humans. This is because text-text semantic comparison is much easier than cross-modal seman-

Table 2: Ablation Study on Flickr8k-Expert (F-Ex), Flickr8k-CF (F-CF), Composite (Com), Pascal-50S and THumB 1.0. HTN denotes using hard text negatives in coarse-grained score learning, and FS refers to the fine-grained score learning. ‘ $v_g$ ’ means incorporating the vision global feature in the Intra&Inter Modality Fusion Module.

Id	Architecture			Training		F-Ex	F-CF	Com	HC	Pascal-50S				THumB w/o h			THumB w/h		
	Intra	Inter	$v_g$	HTN	FS					HI	HM	MM	Mean	P	R	Total	P	R	Total
r1	✓		✓			51.7	36.8	57.8	58.0	99.5	95.0	76.3	82.2	0.23	0.26	0.35	0.20	0.26	0.32
r2		✓	✓			55.1	<b>37.1</b>	59.0	59.5	99.8	95.4	78.1	83.2	0.23	0.26	0.35	0.20	0.26	0.32
r3	✓	✓				55.1	36.9	<b>59.4</b>	58.6	<b>99.9</b>	95.7	79.6	83.5	0.21	0.26	0.34	0.19	0.26	0.32
r4	✓	✓	✓			<b>55.2</b>	36.9	59.3	58.0	99.7	96.1	<b>80.8</b>	83.7	0.22	0.26	0.35	0.20	0.26	0.33
r5	✓	✓	✓	✓		54.5	36.2	58.8	<b>69.3</b>	99.6	93.7	75.2	84.5	<b>0.23</b>	0.28	<b>0.37</b>	<b>0.22</b>	0.30	0.37
r6	✓	✓	✓		✓	<b>55.2</b>	37.0	59.3	60.2	99.7	96.8	79.6	84.1	0.22	0.26	0.34	0.20	0.26	0.32
r7	✓	✓	✓	✓	✓	54.2	36.3	59.2	69.0	99.8	94.0	78.3	<b>85.3</b>	0.22	<b>0.30</b>	<b>0.37</b>	0.21	<b>0.32</b>	<b>0.38</b>

Table 3: Experiments on THumB 1.0. ‘w/o Human’ means discarding human annotated image-caption pairs.

Ref Metric	w/o Human			w/ Human		
	P	R	Total	P	R	Total
w/ BLEU	.21	.13	.25	.15	.04	.13
ROUGE-L	.26	.17	.31	.18	.07	.18
CIDEr	.27	.18	.33	.21	.11	.23
SPICE	.26	.15	.30	.20	.09	.21
BERT-S	.27	.18	.33	.20	.10	.21
RefCLIP-S	<b>.34</b>	<b>.27</b>	<b>.44</b>	<b>.31</b>	<b>.26</b>	<b>.41</b>
w/o InfoCLIP <sup>R</sup>	.05	.19	.17	.05	.19	.17
InfoCLIP <sup>P</sup>	.11	-.22	-.08	.09	-.20	-.08
InfoCLIP	.13	-.06	.04	.11	.06	.03
InfoCLIP <sub>tune</sub>	.15	-.15	.00	.11	-.15	-.03
CLIP-S	.18	.27	.32	.17	.28	.32
CLIP-S <sub>tune</sub>	.15	.26	.29	.13	.26	.28
InfoMetIC <sup>R</sup>	.18	.29	.34	.19	.32	.36
InfoMetIC <sup>P</sup>	.23	.27	.36	.20	.27	.33
InfoMetIC	.22	.30	.37	.21	.32	.38
InfoMetIC <sup>+</sup>	<b>.22</b>	<b>.33</b>	<b>.39</b>	<b>.21</b>	<b>.34</b>	<b>.39</b>

tic comparison, making the precision correlation of reference-based metrics higher. However, limited textual references cannot fully capture image contents, which is harmful for vision recall. **Finally**, InfoMetIC achieves much better performance than InfoCLIP, which shows the effectiveness of our proposed modules on top of CLIP.

#### 4.2.4 Ablation Study

We first validate the effectiveness of our model architecture. As shown in Table 2, removing Intra-modal encoders (r2 vs r4) or Inter-modal encoder (r1 vs r4) results in performance drop on Flickr8k-Expert, Composite and Pascal-50S. Besides, removing global vision feature  $v_g$  from Intra&Inter encoding (r3 vs r4) leads to slight performance drop on Flickr8k-Expert, Pascal-50S and THumB 1.0.

We then carry out ablation study to verify the effectiveness of our training strategy in Table 2. Our proposed hard textual negatives (r4 vs r5) achieves

Table 4: Cross-modal retrieval performances on Nocaps.

Method	image to text			text to image		
	R@1	R@5	R@10	R@1	R@5	R@10
TIGER	63.8	87.0	92.4	22.5	66.5	81.9
CLIP-S	88.2	98.3	99.7	67.5	91.5	95.8
InfoMetIC	76.6	96.5	99.1	71.6	94.4	97.7
InfoMetIC <sup>+</sup>	<b>90.9</b>	<b>98.8</b>	<b>99.7</b>	<b>76.2</b>	<b>95.9</b>	<b>98.4</b>

significant improvements on HC subset of Pascal50s and THumB 1.0 Recall. This shows that constructing hard negatives indeed helps model better evaluate the vision content recall. Adding fine-grained score learning task (r4 vs r6) is also beneficial to the performance of coarse-grained score, which performs better on Pascal-50S and is comparable on other datasets. When trained with all tasks together (r7), InfoMetIC further improves on Pascal-50S and THumB 1.0, and achieves state-of-the-art performance on all datasets.

#### 4.3 Generalization Ability

InfoMetIC are trained with image-captions of Flickr30k and MSCOCO. To evaluate its generalization ability, we further conduct experiments on NoCaps (Agrawal et al., 2019), whose objects are greatly different from Flickr30k and MSCOCO. Since there are no human-labeled scores for image-caption pairs, we perform text-image cross-modal retrieval to validate the effectiveness of our metric. As shown in Table 4, InfoMetIC performs worse than CLIP-S on image-to-text retrieval but better on text-to-image retrieval. After combining with CLIP similarity, InfoMetIC<sup>+</sup> achieves the state-of-the-art performance on both two retrieval tasks. It indicates our overall score can also perform well on instances with unseen objects.

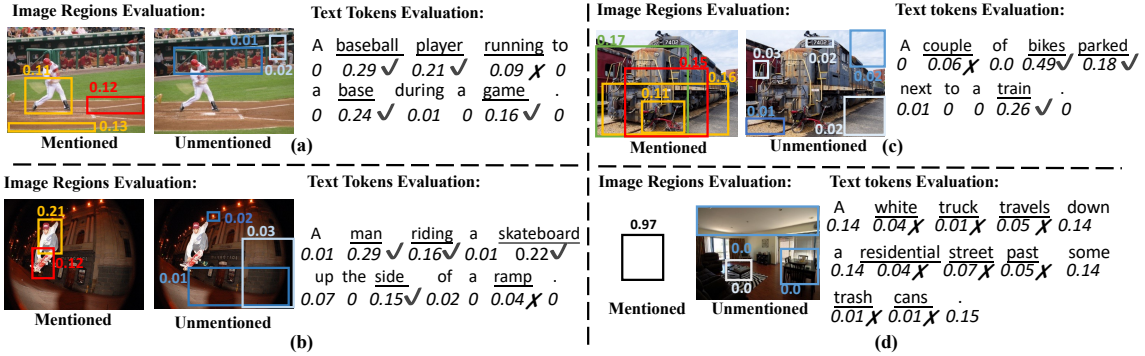


Figure 3: Fine-grained evaluation examples using InfoMetIC. Semantic words in captions are underlined. Semantic words and image regions with a score lower than 0.1 are considered as incorrect words and unmentioned regions respectively. We show all mentioned regions but only the top unmentioned image regions for better visualization.

#### 4.4 Fine-grained Score Evaluation

**Dataset.** To validate the token-level evaluation performance of InfoMetIC, we collect a fine-grained caption evaluation benchmark called CapTokenEval. CapTokenEval is built upon a subset of THumB 1.0. We select 700 image-caption pairs whose precision scores are not perfect ( $< 5.0$ ). For the text part, annotators are asked to judge which words are irrelevant with the image. For the image part, we collect 20 bounding boxes and ask annotators to identify mentioned regions. More details about the annotation can be found in Appendix E.

**Quantitative Results.** Given each image-caption pair, InfoMetIC produces sequence of prediction for both image regions and caption tokens. To quantify token-level evaluation performance, for the text part, we only calculate the accuracy of semantic tokens (nouns, verbs, adjectives and numbers). As shown in Table 5, without extra parameters, InfoCLIP achieves promising performance for fine-grained visual evaluation but poor performance in the text part. Consistent with the result shown in Table 3 that InfoCLIP<sup>R</sup> outperforms InfoCLIP<sup>P</sup>, it further shows the importance of context fusion for text precision evaluation. With multi-task learning, InfoMetIC achieves promising prediction accuracy on both vision and text sequence. Both hard textual negatives and fine-grained score learning task contribute to token-level evaluation performance. Notably, fine-grained score learning task greatly boosts the text-part accuracy. Coarse-grained contrastive learning for text precision score within a batch can result in the model only putting relatively higher weights on a few correct text tokens. Our fine-grained score learning task could effectively

Table 5: Token-level evaluation on CapTokenEval. CS: coarse-grained score learning; HTN: adding hard textual negatives in CS; FS: fine-grained score learning.

Method	Training			Accuracy	
	CS	HTN	FS	Vision	Text
InfoCLIP	-	-	-	0.73	0.33
InfoCLIP <sub>tune</sub>	-	-	-	0.74	0.37
Ours	✓	×	×	0.74	0.36
	✓	✓	×	0.75	0.37
	✓	×	✓	0.75	0.79
	✓	✓	✓	<b>0.75</b>	<b>0.80</b>

alleviate this lazy behavior by teaching the model to put high weights on all correct tokens.

**Qualitative Results.** We show some qualitative results of token-level evaluation in Figure 3. Firstly, InfoMetIC is able to identify various mistakes made in captions, including wrong actions (e.g. “running” in case a), wrong objects (e.g. “ramp” in case b), and wrong modifiers (e.g. “couple” in case c). Secondly, InfoMetIC could report mentioned image regions (e.g. the “skateboard” region in case b) and unmentioned regions (e.g. the “building” region in case b). Especially, when the caption is totally irrelevant with the image, as shown in case d, InfoMetIC could not only judge the wrong semantic words but also inform that all image regions are not mentioned by putting a very high score to the vision null token. One limitation of current metric is that although we perform region filtering by clustering, we still find some similar regions as shown in Figure 3(c). Better ways to de-duplicate image regions could bring further improvement.



## 5 Conclusion

To provide feedbacks on detailed mistakes of image captions, we propose a reference-free informative metric InfoMetIC based on a state-of-the-art vision-language model. InfoMetIC not only points out incorrect descriptions, but also tells which regions are not mentioned. Based on these fine-grained evaluation, InfoMetIC derives a text precision score, a vision recall score, and an overall score. We design both coarse- and fine-grained training tasks to optimize our metric. The overall score given by our metric achieves state-of-the-art correlation with human judgement on multiple benchmarks. We further build a token-level caption evaluation benchmark CapTokenEval to prove the effectiveness of our fine-grained evaluation.

## Limitations

This work focuses on informative image captioning evaluation, including an overall score, vision recall, text precision and token-level scores. The effectiveness of our metric is validated on standard image captioning benchmarks. InfoMetIC in this work may not perform well in other captioning tasks due to domain gap, but we contend that our general framework can be adapted to other domains such as text-aware image captioning. For example, for text-aware image captioning which focuses more on scene texts in images, we could further encode text regions besides the existing object regions for better comparison with captions. In the future, we will comprehensively explore how to adapt our metric to other captioning tasks, such as text-aware image captioning and video captioning.

## Acknowledgements

This work was partially supported by the National Key R&D Program of China (No.2020AAA0108600) and the National Natural Science Foundation of China (No. 62072462).

## References

Somak Aditya, Yezhou Yang, Chitta Baral, Yiannis Aloimonos, and Cornelia Fermüller. 2018. Image understanding using vision and reasoning through scene description graph. *Comput. Vis. Image Underst.*, 173:33–45.

Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi

Parikh, Stefan Lee, and Peter Anderson. 2019. no-caps: novel object captioning at scale. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8948–8957.

Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. Spice: Semantic propositional image caption evaluation. In *European conference on computer vision*, pages 382–398. Springer.

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 *CVPR*, pages 6077–6086. Computer Vision Foundation / IEEE Computer Society.

Satanjeev Banerjee and Alon Lavie. 2005. METEOR: an automatic metric for MT evaluation with improved correlation with human judgments. In *IJEEvaluation@ACL*, pages 65–72. Association for Computational Linguistics.

Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. 2020. Uniter: Universal image-text representation learning. In *European conference on computer vision*, pages 104–120. Springer.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*, pages 4171–4186. Association for Computational Linguistics.

Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. 2021. Clipscore: A reference-free evaluation metric for image captioning. In *EMNLP (1)*, pages 7514–7528. Association for Computational Linguistics.

Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013a. Framing image description as a ranking task: Data, models and evaluation metrics. *J. Artif. Intell. Res.*, 47:853–899.

Micah Hodosh, Peter Young, and Julia Hockenmaier. 2013b. Framing image description as a ranking task: Data, models and evaluation metrics. *Journal of Artificial Intelligence Research*, 47:853–899.

Anwen Hu, Shizhe Chen, and Qin Jin. 2020. Icecap: Information concentrated entity-aware image captioning. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 4217–4225.

Anwen Hu, Shizhe Chen, and Qin Jin. 2021. Question-controlled text-aware image captioning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pages 3097–3105.

Ming Jiang, Junjie Hu, Qiuyuan Huang, Lei Zhang, Jana Diesner, and Jianfeng Gao. 2019a. Reo-relevance, extraneous, omission: A fine-grained evaluation for image captioning. In *EMNLP/IJCNLP (1)*, pages 1475–1480. Association for Computational Linguistics.

- Ming Jiang, Qiuyuan Huang, Lei Zhang, Xin Wang, Pengchuan Zhang, Zhe Gan, Jana Diesner, and Jianfeng Gao. 2019b. Tiger: Text-to-image grounding for image caption evaluation. In *EMNLP/IJCNLP (1)*, pages 2141–2152. Association for Computational Linguistics.
- Jungo Kasai, Keisuke Sakaguchi, Lavinia Dunagan, Jacob Morrison, Ronan Le Bras, Yejin Choi, and Noah A. Smith. 2022. Transparent human evaluation for image captioning. In *NAACL-HLT*, pages 3464–3478. Association for Computational Linguistics.
- Hwanhee Lee, Seunghyun Yoon, Franck Dernoncourt, Trung Bui, and Kyomin Jung. 2021. UMIC: an unreference metric for image captioning via contrastive learning. In *ACL/IJCNLP (2)*, pages 220–226. Association for Computational Linguistics.
- Hwanhee Lee, Seunghyun Yoon, Franck Dernoncourt, Doo Soon Kim, Trung Bui, and Kyomin Jung. 2020. Vilbertscore: Evaluating image caption using vision-and-language bert. In *Proceedings of the First Workshop on Evaluation and Comparison of NLP Systems*, pages 34–39.
- Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. 2018. Stacked cross attention for image-text matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 201–216.
- Chenliang Li, Haiyang Xu, Junfeng Tian, Wei Wang, Ming Yan, Bin Bi, Jiabo Ye, He Chen, Guohai Xu, Zheng Cao, Ji Zhang, Songfang Huang, Fei Huang, Jingren Zhou, and Luo Si. 2022a. mplug: Effective and efficient vision-language learning by cross-modal skip-connections. In *EMNLP*, pages 7241–7259. Association for Computational Linguistics.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022b. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- 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 *European conference on computer vision*, pages 740–755. Springer.
- Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *NeurIPS*, pages 13–23.
- Yingwei Pan, Ting Yao, Yehao Li, and Tao Mei. 2020. X-linear attention networks for image captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10971–10980.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing 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.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. 2015a. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4566–4575.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015b. Cider: Consensus-based image description evaluation. In *CVPR*, pages 4566–4575. IEEE Computer Society.
- Sijin Wang, Ziwei Yao, Ruiping Wang, Zhongqin Wu, and Xilin Chen. 2021. Faier: Fidelity and adequacy ensured image caption evaluation. In *CVPR*, pages 14050–14059. Computer Vision Foundation / IEEE.
- Qiaolin Xia, Haoyang Huang, Nan Duan, Dongdong Zhang, Lei Ji, Zhifang Sui, Edward Cui, Taroon Bharti, and Ming Zhou. 2021. Xgpt: Cross-modal generative pre-training for image captioning. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 786–797. Springer.
- Haiyang Xu, Ming Yan, Chenliang Li, Bin Bi, Songfang Huang, Wenming Xiao, and Fei Huang. 2021. E2e-vlp: End-to-end vision-language pre-training enhanced by visual learning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 503–513.
- Yanzhi Yi, Hangyu Deng, and Jinglu Hu. 2020. Improving image captioning evaluation by considering inter references variance. In *ACL*, pages 985–994. Association for Computational Linguistics.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2:67–78.

Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. Vinvl: Revisiting visual representations in vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with BERT. In *ICLR*. OpenReview.net.

Luowei Zhou, Hamid Palangi, Lei Zhang, Houdong Hu, Jason J. Corso, and Jianfeng Gao. 2020. Unified vision-language pre-training for image captioning and VQA. In *AAAI*, pages 13041–13049. AAAI Press.

## A Context Overuse Issue

CLIP (Radford et al., 2021) is trained to well align global image representations and sentence representation. Thus it applies a triangle masking during text encoding and treats the representation of the last text token [e] as the sentence representation. Due to the training objective and text masking mechanism, the text context information is accumulated with the sequence order, which is unfavorable for text-part fine-grained evaluation. As shown in Figure 4, the third ‘a’ is a meaningless indefinite article but gets a higher relevance score than the correct noun ‘man’.

## B Salience of Visual Information

Our vision recall score is calculated by comparing the text-conditioned vision features (the CLIP’s global vision feature) rather than the sum or average of all regions features. CLIP is trained with massive image-caption pairs and achieves promising performance on multiple Vision-Language tasks. Thus it’s convincing that the global vision feature produced by CLIP could well represent the salient information in an image. As illustrated in Figure 5, both ‘cloud’ and ‘grass’ are objects in the image, but InfoMetIC gives the second caption higher vision recall score because ‘grass’ is more salient than ‘clouds’ in the image.

## C Cluster Number Setting Details

Similar image regions can cause confusion during fine-grained evaluation. In this work, redundant regions are removed by K-means clustering algorithm. Concretely, with 100 bounding boxes given by the object detection model, we perform

Table 6: Performance of InfoMetIC with different cluster numbers on Flickr8k-Expert (F-Ex), Flickr8k-CF (F-CF), Composite (Com), Pascal-50S and THumB w/ Human.

cluster	F-Ex	F-CF	Com	Pascal50S	Thumb
10	54.2	36.1	58.3	84.8	0.36
20	54.2	36.3	59.2	85.3	0.38
30	54.4	36.3	59.5	85.2	0.36
40	54.7	36.2	59.2	85.3	0.39
50	54.8	36.3	59.5	85.3	0.37

K-means to generate  $N$  clusters. For each cluster, the region with highest confidence score given by the object detection model is maintained. The evaluation performance of InfoMetIC with different  $N$  settings is shown in Table 6. With the cluster number ranging from 10 to 50, the overall evaluation performance of InfoMetIC shows minor difference on these benchmarks. Taking into account both performance and complexity, we finally set  $N$  as 20.

## D Baseline Details

To verify the effectiveness of InfoMetIC, besides state-of-the-art caption metrics, we set extra three baselines CLIP- $S_{tune}$ , InfoCLIP and InfoCLIP $_{tune}$ . As shown in Figure 6(a), CLIP-S (Hessel et al., 2021) directly uses the global representations given by CLIP (Radford et al., 2021) to calculate a cosine similarity as the overall score. CLIP- $S_{tune}$  follows the same calculation manner but uses a CLIP fine-tuned on MSCOCO and Flickr30k as the backbone. Previous metrics can’t do fine-grained caption evaluation. Therefore, we set a fine-grained evaluation baseline InfoCLIP, as shown in Figure 6(b). InfoCLIP performs fine-grained scoring as InfoMetIC without Intra&Inter Modality Fusion and parameters in Fine-grained Scoring, e.g.  $W_q^v$  and  $W_k^v$  in Eq (1). InfoCLIP $_{tune}$  means using a fine-tuned CLIP as the backbone.

## E CapTokenEval Annotation Details

To quantify caption evaluation performance at token level, we collect a fine-grained caption evaluation benchmark called CapTokenEval. The details of our annotation are introduced in following subsections.

### E.1 Data Preparation

We prepare image-caption pairs for annotation based on the publicly released dataset THumB 1.0



Caption: A man with a red helmet on a small moped on a dirt road.

Token [s]	a	man	with	a	red	helmet	on	a	small	mo	ped	on	a	dirt	road	.	[e]	
CLIP-S	18.6	9.96	13.45	17.08	12.4	17.95	16.02	19.3	18.3	17.72	21.2	22.77	19.38	21.25	18.53	23.75	26.1	32.62

Figure 4: An illustration about the context overuse during text encoding of CLIP. The CLIP-S of each token are calculated with global vision feature and token-level text feature got by original CLIP encoding way rather than individually encoding.



Caption	InfoMetIC <sup>R</sup>
A very large sheep is standing.	1.66
A very large sheep is standing in the grass.	3.80
A very large sheep is standing under clouds.	2.60

Figure 5: An illustration about the influence of object saliency to our vision recall score InfoMetIC<sup>R</sup>.

(Kasai et al., 2022). THumB 1.0 collects 500 images from MSCOCO (Lin et al., 2014) and pairs each image with 4 captions generated by state-of-the-art image captioning models, including UP-Down (Anderson et al., 2018), Unified-VLP (Zhou et al., 2020), VinVL-base and VinVL-large (Zhang et al., 2021). There are a precision score, a recall score and a total score for each image-caption pair. To ensure that textual token-level evaluation in our benchmark is hard enough, we select image-caption pairs whose precision score is not perfect (<5.0). We finally collect 700 image-captions pairs from ThumB 1.0. As the data used in our annotation all come from publicly released datasets, there are no ethic issues.

For each image, we extract 100 bounding boxes with pre-trained object detection model Bottom-Up (Anderson et al., 2018). To filter similar image regions, we apply K-means clustering on these bounding boxes. We generate 20 clusters for each image and choose a bounding box with highest confidence score of object classification from each cluster. Thus, for each image-caption pair, we provide 20 image regions to annotators, who will choose which regions are mentioned by the caption. For the text part, we tokenize the caption with Spacy<sup>1</sup>.

<sup>1</sup><https://spacy.io/usage>

## E.2 Annotation Platform

We build a platform to support the fine-grained annotation. Figure 7 presents the annotation interface on our platform, which consists of three major parts. The middle part contains an image-caption pair to be annotated. The left part is the textual token-level annotation area, which lists all tokens in the caption. The right part is the visual token-level annotation area, which places 20 images with bounding boxes indicating different image regions.

## E.3 Annotation Instruction

Given an image-caption pair, we ask annotators to identify which tokens in the caption are incorrect and which regions are mentioned by the caption. Besides, we require that if the caption mentions an object without descriptions about details, the image regions of detailed components shouldn't be classified as 'Mentioned'. For example, for the caption 'a group of people riding on the back of an elephant', the image region of the elephant nose shouldn't be judged as 'Mentioned'.

We invite 20 college students as annotators. They all have sufficient English proficiency to understand image captions in English. We provide a document to inform annotators the goal of our annotation and detailed instructions about the usage of the annotation platform. Each annotator is assigned 35 image-caption pairs for annotation.

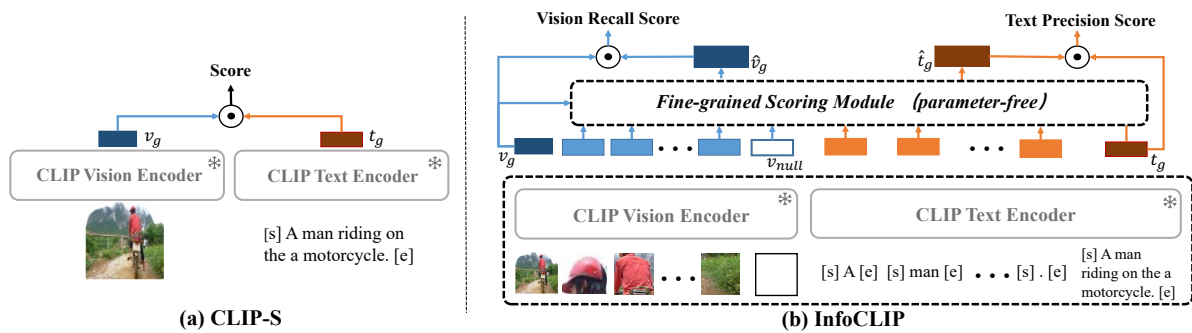


Figure 6: Overall architectures of baselines CLIP-S and InfoCLIP

### CapTokenEval Annotation

**Text Annotation** (which words are correct?)

A @ True  False

group @ True  False

of @ True  False

people @ True  False

riding @ True  False

on @ True  False

the @ True  False

back @ True  False

of @ True  False

an @ True  False


elephant @ True  False

. @ True  False

**Vision Annotation** (which regions are mentioned by the caption?)

0 id 0-34

Go! < >

 COCO\_val014\_00000000974.jpg

A group of people riding on the back of an elephant.





















			
<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False
			
<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False
			
<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False
			
<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False
			
<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False	<input type="radio"/> True <input type="radio"/> False

Figure 7: An overview of the annotation platform.

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*Left blank.*
- A2. Did you discuss any potential risks of your work?  
*Left blank.*
- A3. Do the abstract and introduction summarize the paper’s main claims?  
*Left blank.*
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

*Left blank.*

- B1. Did you cite the creators of artifacts you used?  
*Left blank.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*Left blank.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*Left blank.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*Left blank.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*Left blank.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*Left blank.*

### C Did you run computational experiments?

*Left blank.*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
*Left blank.*

---

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*Left blank.*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*Left blank.*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

*Left blank.*

**D**  **Did you use human annotators (e.g., crowdworkers) or research with human participants?**

*Left blank.*

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*Left blank.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*Left blank.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*Left blank.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*Left blank.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*Left blank.*