

iParaphrasing: Extracting Visually Grounded Paraphrases via an Image

Chenhui Chu,¹ Mayu Otani² and Yuta Nakashima¹

¹Institute for Datability Science, Osaka University

²CyberAgent, Inc.

chu,n-yuta@ids.osaka-u.ac.jp, otani_mayu@cyberagent.co.jp

Abstract

A paraphrase is a restatement of the meaning of a text in other words. Paraphrases have been studied to enhance the performance of many natural language processing tasks. In this paper, we propose a novel task *iParaphrasing* to extract *visually grounded paraphrases (VGPs)*, which are different phrasal expressions describing the same visual concept in an image. These extracted VGPs have the potential to improve language and image multimodal tasks such as visual question answering and image captioning. How to model the similarity between VGPs is the key of *iParaphrasing*. We apply various existing methods as well as propose a novel neural network-based method with image attention, and report the results of the first attempt toward *iParaphrasing*.

1 Introduction

A paraphrase is a restatement of the meaning of a word, phrase, or sentence within the context of a specific language (e.g., “a red jersey” and “a red uniform shirt” in Figure 1 are paraphrases) (Bhagat and Hovy, 2013). Paraphrases have been exploited for natural language understanding, and shown to be very effective for various natural language processing (NLP) tasks, including question answering (Riezler et al., 2007), summarization (Zhou et al., 2006), machine translation (Chu and Kurohashi, 2016), text normalization (Ling et al., 2013), textual entailment recognition (Androutsopoulos and Malakasiotis, 2010), and semantic parsing (Berant and Liang, 2014).

In this paper, we propose a novel task named, *iParaphrasing*, to extract *visually grounded paraphrases (VGPs)*. We define VGPs as different phrasal expressions that describe the same visual concept in an image. Nowadays, with the spread of the web and social media, it is easy to collect large amounts of images with their describing text. For example, different news sites release news with the same topic using the same image; photos with many comments are posted to social networking sites and blogs. As these describing texts are written by different people but about the same image, there are potentially large amounts of VGPs in the describing text (Figure 1). We aim to accurately extract these paraphrases using the image as a pivot to associate different phrases.

The extracted VGPs can be applied to various computer vision (CV) and NLP tasks, such as image captioning (Vinyals et al., 2015) and visual question answering (VQA) (Wu et al., 2017), for the better understanding of both images and languages. For example, a VQA system must understand queries of different expressions about the same visual concept (e.g., “a male” and “the pitcher” in Figure 1) in order to answer a question properly. VGPs can also be applied to the evaluation of image captioning systems in the similar way as paraphrases have been applied for machine translation evaluation (Snover et al., 2009).

As a pioneering study, we work on *iParaphrasing* on the Flickr30k entities dataset (Plummer et al., 2015). This dataset contains 30k images with 5 captions per image annotated via crowdsourcing, which can be seen as a very small subset of the data available in the web and social media. Figure 1 shows an example image together with its five captions taken from this dataset. In the Flickr30k entities dataset, entities (i.e., noun phrases) in the captions have been manually aligned to their corresponding image

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

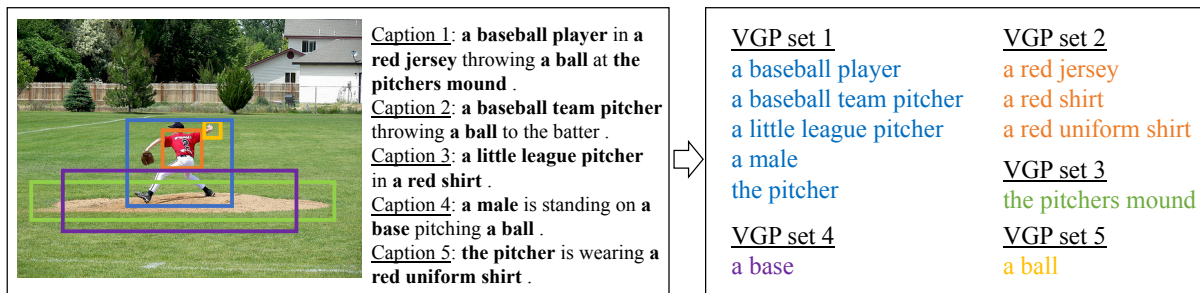


Figure 1: An example from the Flickr30k entities dataset, in which an image is described by five captions (entities in the captions are marked in bold). Our task is to extract the entities that describe the same visual concept (represented as an image region) in the image as VGPs. Note that the image regions are not given as input but are drawn here for comprehensibility.

regions (Plummer et al., 2015). Therefore, we can obtain a set of phrases annotated with the same image region. This set of phrases are used as the ground truth VGPs in our study. The goal of this work is to extract these VGPs.

We formulate our task as a clustering task (Section 3), where the similarity between each entity pair is crucial for the performance. We apply many different unsupervised similarity computation methods (Section 4) including phrase localization-based similarity (Plummer et al., 2017) (Section 4.1), translation probability-based similarity (Koehn et al., 2007) (Section 4.2), and embedding-based similarity (Mikolov et al., 2013; Klein et al., 2014; Plummer et al., 2015) (Section 4.3). In addition, we propose a supervised neural network (NN)-based method using both textual and visual features to explicitly model the similarity of an entity pair as VGPs (Section 5). Experiments show that our proposed NN-based method outperforms the other methods.¹

2 Related Work

2.1 Paraphrase Extraction

Previous studies extract paraphrases from either monolingual corpora or bilingual parallel corpora. One major approach is to use the distributional similarity (Harris, 1954) with regular monolingual corpora (a large collection of text in a single language) (Lin and Pantel, 2001; Bhagat and Ravichandran, 2008; Marton et al., 2009), or monolingual comparable corpora (a set of monolingual corpora that describe roughly the same topic in the same language) (Barzilay and Lee, 2003; Chen and Dolan, 2011). Distributional similarity stems from the distributional hypothesis (Harris, 1954), stating that words/phrases that share similar meanings should appear in similar distributions. This approach sometimes suffers from noisy results, because the distributed similarity often maps antonyms to closer points. Some methods try to extract paraphrases from monolingual parallel corpora (a collection of sentence level paraphrases) (Arase and Tsujii, 2017; MacCartney et al., 2008), but such monolingual parallel corpora are rarely available.

Bilingual parallel corpora (a collection of sentence-aligned bilingual text) enjoys more availability than monolingual parallel corpora as they are mandatory for training machine translation systems. Bilingual parallel corpora can be used for paraphrase extraction, with bilingual pivoting (Bannard and Callison-Burch, 2005). This method assumes that two source phrases are a paraphrase pair if they are translated to the same target phrase. Bilingual pivoting has been further refined by using syntax information (Callison-Burch, 2008) or mutual information (Kajiwara et al., 2017). These methods have led to the construction of a multilingual paraphrase database (Ganitkevitch and Callison-Burch, 2014).

Note that our definition of paraphrases may look different from the studies mentioned above, as our paraphrases are a set of noun phrases that represent the same visual concept. Our idea to extract paraphrases under this definition is to use image captioning datasets (Young et al., 2014; Chen et al., 2015),

¹Codes and data for reproducing the results reported in this paper are available at https://github.com/ids-cv/coling_iparaphrasing

which usually contain several captions for each image, and currently scale to sub-million images, instead of a bilingual parallel corpus with limited availability. To the best of our knowledge, this is the first study that aims to extract paraphrases from such multimodal datasets consisting of images and their captions.²

2.2 Coreference Resolution

Coreference resolution is a task to find the expressions that refer to the same entity in a text (Soon et al., 2001; Lee et al., 2017). Our task in this paper focuses on extracting entities that describe the same visual concept, making the formulation similar to coreference resolution. Our task differs from conventional coreference resolution that it requires visual grounding. In addition, the targets of coreference resolution are the entities in a sentence or a document, while our targets are the entities in the captions of an image that are quasi-paraphrases but are not related to each other in discourse level like sentences in a document. For coreference resolution, the context in a sentence or discourse information in a document are crucial, but discourse information does not exist in our task.³ In the context of vision and language tasks, Konget al. (2014) used noun/pronoun coreference resolution in sentential descriptions of RGB-D scenes for improving 3D semantic parsing. The texts being handled in their work are either sentences or documents, and their targets are limited to noun words and pronouns but we extract noun phrases. Because our goal is not limited to entities but arbitrary phrases, we believe that comparing to coreference resolution, iParaphrasing is a more forethoughtful name to define our task for future research.

2.3 Phrase Localization

Phrase localization is a task to find an image region that corresponds to a given phrase in a caption, which is closely related to our VGP extraction task. Plummer et al. (2015) pioneered this work, in which they annotated phrase-region alignment in the Flickr30k image-caption dataset (Young et al., 2014) and released it as the Flickr30k entities dataset. They also proposed a method based on canonical correlation analysis (CCA) (Hardoon et al., 2004) that learns joint embeddings of phrases and image regions for associating them. Wang et al. (2016a) proposed joint embeddings using a two-branch NN. Fukui et al. (2016) used a multimodal compact bilinear pooling method to combine textual and visual embeddings. Rohrbach et al. (2016) proposed a convolutional NN (CNN)-recurrent NN (RNN)-based method for this task. They learn to detect a region for a given phrase and then reconstruct the phrase using the detected region. Wang et al. (2016b) noticed that the relationships between phrases should agree with their corresponding regions, and proposed a joint matching method, but their method only considers the “has-a” relationship that is explicitly indicated by possessive pronouns. Previous studies rely on region proposal to produce a number of region candidates for phrase localization, Yeh et al. (2017) proposed a unified framework that can search over all possible regions. Plummer et al. (2017) used spatial relationships between pairs of entities connected by verbs or prepositions, which achieved the state-of-the-art performance. In this paper, we use the current state-of-the-art phrase localization method of (Plummer et al., 2017) as a baseline for VGP extraction.

2.4 Other Vision and Language Tasks

Vision and language tasks have been a hot research area recently in both the CV and NLP communities. Various efforts have been made for many multimodal tasks such as image object/region referring expression grounding (Kazemzadeh et al., 2014; Mao et al., 2016; Hu et al., 2017; Cirik et al., 2018) visual captioning (Vinyals et al., 2015; Xu et al., 2015; Bernardi et al., 2016; Laokulrat et al., 2016), text-image retrieval (Otani et al., 2016), visual question answering (Wu et al., 2017), visual dialog (Das et al., 2017a; Das et al., 2017b) and video event detection (Phan et al., 2016). Some researchers also have employed images for improving NLP tasks, such as multimodal machine translation (Specia et al., 2016), cross-lingual document retrieval (Funaki and Nakayama, 2015), and textual entailment recognition (Han

²Although (Plummer et al., 2015) annotated the VGPs in the Flickr30k entities dataset, they did not propose any methods to extract them. (Regneri et al., 2013) collected sentence level paraphrases by aligning video scripts with the same time frame; these sentence level paraphrases are essentially similar to captions of an image.

³If we treat multiple captions as a document forcibly, we could apply a coreference resolution approach for our task.




Figure 2: An overview of our VGP extraction formulation. We extract VGP via clustering, where the entity-entity similarity $s(i, j)$ is the key. We compare both unsupervised and supervised methods using entity-image and entity-entity associations for computing this similarity.

et al., 2017). iParaphrasing is a novel CV+NLP task, which to the best of our knowledge has not been studied before and can boost the performance of various multimodal and NLP tasks.

3 Paraphrase Extraction via Clustering

We formulate the paraphrase extraction from the Flickr30k entities dataset as a clustering task. Given an image and all the entities in the corresponding captions, the task is to cluster the entities⁴ to its corresponding visual concepts represented as image regions. The number of clusters (i.e., the number of paraphrase sets in a set of an image and captions) is not explicitly given in our task. Therefore, we apply the affinity propagation algorithm (Frey and Dueck, 2007) to cluster entities, which can estimate the number of clusters as well.

Affinity propagation creates clusters by iteratively sending two types of messages between pairs of entities until convergence. The first type is the responsibility $r(i, j)$ sent from entity i to candidate representative entity j , indicating the strength that entity j should be the representative entity for entity i , which is defined as:

$$r(i, j) \leftarrow s(i, j) - \max_{\forall j' \neq j} \{a(i, j') + s(i, j')\} \quad (1)$$

where $s(i, j)$ is the similarity between entities i and j . The second type is the availability $a(i, j)$ sent from candidate representative entity j to entity i , indicating to what degree that candidate representative entity j is the cluster center for entity i , which is defined as:

$$a(i, j) \leftarrow \min \left\{ 0, r(j, j) + \sum_{\forall i' \notin \{i, j\}} \max \{0, r(i', j)\} \right\} \quad (2)$$

At the beginning, the values of $r(i, j)$ and $a(i, j)$ are set to zero, and they are updated in every iteration until convergence. We optimize the number of clusters on a validation split by adjusting the preference (i.e., self similarity $s(i, i)$) of affinity propagation.

Figure 2 shows an overview of our formulation, where the similarity between the entities is the key. We apply various unsupervised methods for computing this similarity, and propose a supervised NN-based model.

4 Unsupervised Similarity Methods

We apply phrase localization for modeling the entity-entity similarity based on entity-image association (Section 4.1). In addition, we apply various methods for modeling the entity-entity similarity directly (Sections 4.2, and 4.3).

⁴In this paper, we assume that entities are given. In the case that entities are not given, we can easily extract them by chunking the noun phrases.

4.1 Phrase Localization-Based Similarity

The similarity between entities i and j is defined as:

$$s(i, j) = \sum_{r_m \in R} p(i|r_m)p(j|r_m) \quad (3)$$

where R is a set of image regions that are aligned to both entities i and j obtained with the phrase localization method of (Plummer et al., 2017); $p(i|r_m)$ is the localization probability of r_m for i , defined as:

$$p(i|r_m) = \frac{l(i, r_m)}{\sum_{r_m \in R} l(i, r_m)} \quad (4)$$

where $l(i, r_m)$ is the localization score of region r_m for entity i obtained using the method of (Plummer et al., 2017).

4.2 Translation Probability-Based Similarity

The similarity between entities i and j is defined as:

$$s(i, j) = p(i|j)p(j|i) \quad (5)$$

where $p(i|j)$ and $p(j|i)$ are the direct and inverse translation probabilities of an entity pair i and j , which are calculated using a conventional statistical machine translation (SMT) (Koehn et al., 2007) method:

1. Generate a pseudo parallel corpus using the captions in the dataset, which treats the 5 captions for each image as monolingual parallel sentences and pair each of the sentences that leads to $\binom{5}{2} = 10$ sentence pairs per image.
2. Apply word alignment to the parallel corpus using IBM alignment models (Brown et al., 1993) in two directions with the grow-diag-final-and heuristic (Koehn et al., 2007) to align the words in each caption pair.
3. From the word-aligned parallel corpus, extract entity pairs such that the words inside an entity pair are aligned. Then $p(i|j)$ and $p(j|i)$ are calculated as follows:

$$p(i|j) = \frac{c(i, j)}{\sum_k c(i, k)}, \quad p(j|i) = \frac{c(i, j)}{\sum_k c(j, k)} \quad (6)$$

where $c(i, j)$ is the number of co-occurrence of i and j in the word-aligned corpus.

4.3 Embedding-Based Similarity

In this method, the similarity between entities i and j is defined as:

$$s(i, j) = \frac{\mathbf{t}_i^\top \mathbf{t}_j}{\|\mathbf{t}_i\| \|\mathbf{t}_j\|} \quad (7)$$

where \mathbf{t}_i and \mathbf{t}_j are the phrase embeddings of i and j . We compare three different methods for phrase embeddings.

4.3.1 Word Embedding Average

We represent each word with a 300 dimensional word2vec (Mikolov et al., 2013) vector pre-trained on the Google News corpus.⁵ We remove stop words in each entity, and calculate the representation of each entity using the average of all word embeddings.

⁵<https://github.com/mmihaltz/word2vec-GoogleNews-vectors>

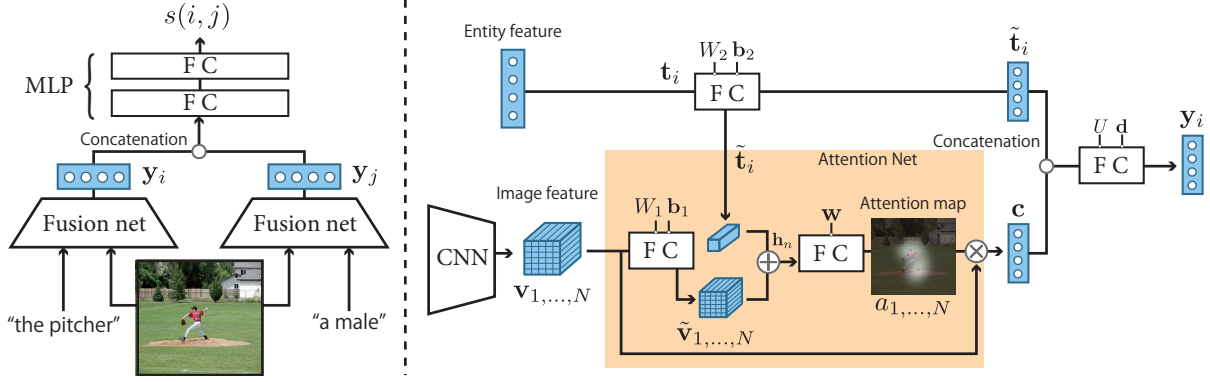


Figure 3: Our supervised NN with image attention-based similarity model (left) and its fusion sub-network (right).

4.3.2 Fisher Vector

Fisher vector is a pooling over word2vec vectors of individual words (Klein et al., 2014), which has been used in the phrase localization task for representing the entities (Plummer et al., 2015). To compute the Fisher vector for an entity, we represent the entity by the HGLMM Fisher vector encoding (Klein et al., 2014) of the word vectors, following (Plummer et al., 2015).⁶

4.3.3 Fisher Vector with CCA

Projecting the feature vectors of image regions and entities to a shared semantic space can provide strong associations between the image regions and entities, which has the potential to improve the performance of VGP extraction. Therefore, we learn a CCA projection on the Flickr30k entities dataset for the image region feature vectors and entity feature vectors with (Plummer et al., 2015), in which the normalized CCA formulation of (Gong et al., 2014) is used. The columns of the CCA projection matrices are scaled by the eigenvalues, and the feature vectors are projected by these matrices and normalized to the dimensionality of 4,096. The image region feature vectors are extracted using Faster R-CNN (Ren et al., 2015).⁷ We use Fisher vectors for entity feature vectors.

5 Supervised Similarity Model Based on Neural Network with Image Attention

We propose a NN-based supervised model. This model computes the similarities of entity pairs as VGPs by explicitly modeling the associations between them and an image. Figure 3 illustrates our proposed NN model.⁸ Given an entity pair and its corresponding image, we construct two separated *fusion nets* for each entity (Figure 3 (right)). Note that parameters of these two fusion net are shared. A fusion net represents an entity with a concatenation of its entity feature vector and visual context vector. The visual context vector is computed with an attention mechanism, indicating to which part of the image should be paid attention, in order to judge whether the entity pair is VGP or not. The outputs of the two fusion nets are then fed into a multilayer perceptron (MLP) to compute the similarity of the two entities.

Formally, let V be a 196×512 feature map⁹ extracted from the `conv5_3` layer in the VGG-16 network (Simonyan and Zisserman, 2015) for an input image; \mathbf{v}_n is a 512 dimensional vector at position n of V . Given an entity feature vector \mathbf{t}_i and \mathbf{v}_n , we first transform them with fully connected (FC) layers whose unit sizes are 512:

$$\tilde{\mathbf{v}}_n = \text{norm}_{L_2}(W_1 \mathbf{v}_n + \mathbf{b}_1), \quad \tilde{\mathbf{t}}_i = \text{norm}_{L_2}(W_2 \mathbf{t}_i + \mathbf{b}_2) \quad (8)$$

⁶The Fisher vector is constructed with 30 centers of both first and second order information, which results in a very sparse vector whose dimensionality is $30 \times 30 \times 2 = 18000$. Therefore, we apply principal component analysis (PCA) to convert it to a lower dimensionality of 4,096.

⁷https://github.com/ShaoqingRen/faster_rcnn

⁸(Yin et al., 2016) proposed a CNN network with attention for sentence level paraphrase identification; our model differs from theirs that we fuse both textual and visual information while theirs is a text only model.

⁹An image is split into $14 \times 14 = 196$ sub-images, and represented as a 196×512 feature map.

where $\text{norm}_{L2}(\cdot)$ indicates L2 normalization to an input vector. We then compute an attention value a_n for \mathbf{v}_n as:

$$\mathbf{h}_n = \text{relu}(\tilde{\mathbf{v}}_n + \tilde{\mathbf{t}}_i), \quad e_n = \mathbf{w}^\top \mathbf{h}_n, \quad a_n = \frac{\exp(e_n)}{\sum_{n=1}^N \exp(e_n)} \quad (9)$$

where $N = 196$. After obtaining a_n , we fuse a visual and an entity feature vector to \mathbf{y}_i as:

$$\mathbf{c} = \sum_{n=1}^N a_n \mathbf{v}_n, \quad \mathbf{y}_i = U[\text{norm}_{L2}(\mathbf{c}), \tilde{\mathbf{t}}_i] + \mathbf{d} \quad (10)$$

where $[\cdot, \cdot]$ indicates the concatenation of two vectors, \mathbf{c} is a visual context vector. We compute fusion feature vectors \mathbf{y}_i and \mathbf{y}_j with the corresponding image. Finally, we feed them to a two-layer MLP network with ReLU non-linearity, whose unit sizes are 128 and 1, respectively, to produce the similarity of the entity pair.

6 Experiments

6.1 Settings

We conducted experiments on the Flickr30k entities dataset (Plummer et al., 2015). This dataset contains 31,837 images, which is described with 5 captions annotated via crowdsourcing. We followed the 29,873 training, 1,000 validation, and 1,000 test image splits used in the phrase localization task (Plummer et al., 2015). Our task is to automatically cluster the entities in the captions that describe the same visual concept (i.e., region in the dataset) in the image as VGPs. Entities that share the same ID and group type (e.g., “a red jersey,” “a red shirt” and “a red uniform shirt” in Figure 1 share the same entity ID and group type “/EN#19026/clothing”) are treated as the ground truth VGP clusters in our evaluation.¹⁰ As stop words should not be considered for computing the entity similarities, we preprocessed the entities in the dataset by removing stop words for all the methods.

We evaluated both clustering and pairwise performance.¹¹ The entity clustering performance for each image was measured with adjusted Rand index (ARI) (Hubert and Arabie, 1985). We used the implementation in the Scikit-learn machine learning toolkit (Thirion et al., 2011)¹² for computing ARI. We report the mean of ARI scores for all the images in the test split. To evaluate the performance for clustering, we optimized the number of clusters by adjusting the preference for affinity propagation on the validation split to maximize the ARI using the Bayesian optimization algorithm (Mockus, 1989) implemented in GPyOpt.¹³ The pairwise performance was evaluated with precision, recall, and F-score, defined as:

$$\text{precision} = \frac{\#\text{predicted_positive}}{\#\text{predicted}}, \quad \text{recall} = \frac{\#\text{predicted_positive}}{\#\text{true_positive}}, \quad \text{F-score} = 2 \times \frac{\text{precision} \times \text{recall}}{\text{precision} + \text{recall}} \quad (11)$$

where an entity pair with a similarity higher than a threshold is treated as *predicted*, which is compared against the ground truth to judge whether it is *predicted_positive* or not. We report the performance using the similarity threshold tuned on the validation split that maximizes the F-score.

We used the affinity propagation implementation¹⁴ in Scikit-learn for clustering. We compared the performance of the different similarity methods described in Sections 4 and 5, where the detailed settings for the methods were as follows:

¹⁰There is an entity type named “notvisual” in the dataset (e.g., “the batter” in Figure 1), which means this entity has no corresponding visual regions in the image. In our evaluation, we excluded this “notvisual” type, because all entities that are not visual are annotated with the same entity ID and thus ground truth VGPs for these “notvisual” entities are unavailable in the dataset. There are entity pairs in the dataset that are the same after removing the stop words (e.g., “a man” and “the man”), we treated them as one entity for evaluation. In addition, entities that do not have corresponding regions in the image were excluded from evaluation.

¹¹We did not report phrase localization accuracies for our proposed NN-based supervised model, because attention is different from phrase localization. Instead of determining the best region corresponding to the given phrases, attention provides attention probabilities to the 196 sub-images, which cannot be used for evaluation directly. A soft-accuracy metric could be reported for discussion, but this metric is not directly comparable to previous phrase localization studies.

¹²<http://scikit-learn.org/stable/modules/clustering.html#adjusted-rand-index>

¹³<https://github.com/SheffieldML/GPyOpt>

¹⁴<http://scikit-learn.org/stable/modules/generated/sklearn.cluster.AffinityPropagation.html>

- Phrase localization (PL): we used the pl-clc toolkit,¹⁵ which is an implementation of the localization method of (Plummer et al., 2017). For R in equation 3, we used the top 30 localization candidates for each entity. The localization scores for each entity and region pair obtained with (Plummer et al., 2017) were used to compute the similarity.
- Translation probability (TP): to get the entity translation probabilities, we first applied the GIZA++ toolkit¹⁶ that is an implementation of the IBM alignment models (Brown et al., 1993) on the pseudo parallel corpus, and then a phrase table was extracted and the phrasal translation probabilities were calculated using state-of-the-art SMT toolkit Moses (Koehn et al., 2007).
- Word embedding average (WEA): see the detailed setting in Section 4.3.1.
- Fisher vector (FV): entity feature vectors were computed using the Fisher vector toolkit released by the authors,¹⁷ following the settings described in Section 4.3.2.
- Fisher vector w/ CCA (FV+CCA): image region feature vectors and entity feature vectors were projected into a 4,096 dimensional space CCA trained on the training split of the Flickr30k entity dataset (Section 4.3.3).
- Supervised NN (SNN): to show the effectiveness of the fusion net (Section 5), we compared a supervised NN-based setting that only feeding the entity feature vectors to the MLP (Figure 3 (left)) for paraphrase similarity prediction. This setting only uses entity feature vectors as input for the NN. It was trained on the training split of the Flickr30k entity dataset. We used all the ground truth VGP pairs in the training split as positive instances. During training, we constructed mini-batches with 15% of positive instances and 85% of randomly sampled negative instances. We used Adam for optimization with a mini-batch size of 300 and weight decay of 0.0001. The learning rate was initialized to 0.01, which was halved at every epoch. We used sigmoid cross entropy loss. We terminated training after 5 epochs, where we observed the loss converged on the validation split. For the entity feature vectors, we compared three different settings described above namely: WEA, FV, and FV+CCA.
- SNN+image: this setting is for our proposed supervised NN-based method described in Section 5. We again compared the three different entity feature vectors. We used VGG-16 (Simonyan and Zisserman, 2015) for the image features. The model was trained with the same configuration as the SNN setting.
- Ensemble: the ensemble of the SNN and SNN+image models that takes the average similarity given by both models. The motivation of this setting is to complement these two models to each other.

6.2 Results

Table 1 shows the results of all the different methods. We report the performance based on the entity types to better understand the performance difference of each method, i.e., “all” evaluates on all entities, whereas “single” and “multi” only evaluate on entities with one single token and multiple tokens, respectively, after removing stop words. For the unsupervised methods, we can see that PL does not show good performance. This is due to the low performance of phrase localization.¹⁸ TP shows a fairly high F-score, but a very low ARI score. The reason for this is that the translation probabilities are computed based on word alignment, leading to a similarity score of 0 to the unrelated entity pairs, which is not suitable for affinity propagation. WEA shows relatively good performance that is better than FV. This is because 45.84% of the entities in our task are single word type after removing the stop words, and converting the low dimensional word embedding to high dimensional and sparse Fisher vectors is harmful

¹⁵<https://github.com/BryanPlummer/pl-clc>

¹⁶<http://code.google.com/p/giza-pp>

¹⁷<https://owncloud.cs.tau.ac.il/index.php/s/vb7ys8Xe8J8s8vo>

¹⁸Although (Plummer et al., 2017) is the current state-of-the-art for phrase localization, the accuracy is only 55.85%.

Method	ARI	Precision	Recall	F-score
	all / single / multi	all / single / multi	all / single / multi	all / single / multi
PL	43.23 / 45.92 / 46.35	59.32 / 51.53 / 62.86	63.12 / 47.99 / 74.14	61.16 / 49.70 / 68.04
TP	37.61 / 50.32 / 36.79	66.23 / 63.20 / 82.17	64.20 / 66.10 / 56.31	65.20 / 64.62 / 66.83
WEA	49.82 / 47.16 / 49.58	61.51 / 46.11 / 62.84	71.29 / 67.93 / 78.47	66.04 / 54.93 / 69.79
FV	39.85 / 43.55 / 41.26	63.79 / 40.84 / 67.51	60.87 / 35.41 / 77.03	62.30 / 37.94 / 71.96
FV+CCA	54.91 / 51.01 / 49.30	64.79 / 55.79 / 68.24	82.20 / 75.83 / 84.98	72.46 / 64.28 / 75.69
SNN (WEA)	60.23 / 55.06 / 53.26	77.86 / 83.66 / 74.50	84.58 / 75.16 / 88.96	81.08 / 79.18 / 81.09
SNN+image (WEA)	60.55 / 55.42 / 55.82	79.47 / 81.01 / 77.26	84.56 / 79.35 / 87.06	81.94 / 80.17 / 81.86
Ensemble (WEA)	60.65 / 54.92 / 54.56	80.65 / 78.68 / 77.38	84.79 / 83.14 / 88.85	82.67 / 80.85 / 82.72
SNN (FV)	48.13 / 45.97 / 47.22	64.21 / 45.92 / 66.40	65.93 / 50.89 / 76.51	65.06 / 48.28 / 71.10
SNN+image (FV)	47.90 / 47.39 / 48.31	63.49 / 52.62 / 66.86	68.20 / 55.62 / 78.01	65.76 / 54.08 / 72.01
Ensemble (FV)	49.82 / 48.16 / 48.34	65.48 / 54.87 / 70.51	71.43 / 56.24 / 76.54	68.33 / 55.55 / 73.40
SNN (FV+CCA)	60.56 / 56.35 / 54.06	83.11 / 85.19 / 77.44	82.13 / 79.30 / 87.69	82.62 / 82.14 / 82.25
SNN+image (FV+CCA)	61.17 / 54.86 / 54.14	82.51 / 84.52 / 80.28	84.19 / 81.85 / 86.82	83.34 / 83.16 / 83.43
Ensemble (FV+CCA)	62.35 / 54.98 / 54.84	82.71 / 84.10 / 80.91	85.67 / 83.50 / 87.06	84.16 / 83.80 / 83.87

Table 1: VGP extraction results (“all” evaluates on all entities, “single” and “multi” only evaluate on entities consist of one single token and multiple tokens after removing stop words, respectively; the methods above and below the double line are unsupervised and supervised, respectively).

for these single word entity pairs. However, for the performance of entities containing multiple words, the Fisher vector is better than word embedding average in the perspective of F-score. FV+CCA significantly outperforms FV. This is because it uses visual information in the training split that transforms the entity vectors and visual vectors into the semantic space that is helpful for detecting VGPs.

Regarding the supervised methods, NN-based methods using any entity feature vectors outperforms the methods that uses them in an unsupervised way. The reason for this is that it directly uses the paraphrase supervision in the training split, while the unsupervised methods do not. Using entity representation with better ARI and F-score for the SNN method can achieve better results. Our proposed method (SNN+image) that uses both textual and visual features shows better performance compared to SNN that uses textual features only, indicating that the usage of visual features is helpful for our VGP extraction task.¹⁹ However, the performance improvements are not very large. We discuss the reason for this in detail in Section 6.3.1. The ensemble of SNN and SNN+image further improves the performance, which means that these two models complement each other.

6.3 Discussion

6.3.1 Neural Network w/ and w/o Images

We compared the SNN and SNN+image results, and found that image attention is helpful for identifying people-related paraphrases in about 50% cases, which are difficult to be determined based on the textual information only. In addition, the attention for these people-related paraphrases are well learned. We believe the reason for this is that many entities in the training split are people-related and thus they are well modeled. Figure 4a shows such an example, where the SNN (FV+CCA) model fails to identify these two entities “a group of order men” and “a group of people” as VGPs due to the diverse textual descriptions of the the same visual concept. Our proposed NN+image (FV+CCA) model correctly identifies these VGPs by paying attention to the image region of people in the image. In about 30% cases, the visual information is also helpful for the identification of other types of paraphrases, although the attention is not accurate. Figure 4b shows an example, where the SNN (FV+CCA) model could not identify two entities “a large display of artifacts” and “an art exhibit” as VGPs.

In about 20% cases, visual information could bring negative effects for paraphrase identification. Figure 4c shows an example that “the street window shops” and “a clothing store window” are mistakenly judged as a paraphrase after using the image information while using textual information judges correctly. Although, the attention for these entities refer to the same visual concept in the image, the entities actually refer to different concepts (i.e., “shop” and “window”).

¹⁹The difference between SNN (FV+CCA) and SNN+image (FV+CCA) is whether using visual features for iParaphrasing explicitly or not. SNN (FV+CCA) uses image region features for learning entity features, but it does not use visual features explicitly for iParaphrasing.




Figure 4: Examples comparing SNN (FV+CCA) with SNN+image (FV+CCA) ((a), (b), and (c)); the leftmost images are the original ones, the images in the middle and on the right show attention of the entity pairs on the images, the degree of whiteness indicates the strength of attention, the identification results are shown under the images), and failed examples (d).

6.3.2 Failed Examples

Even the best method, namely Ensemble (FV+CCA), only achieves a ARI of 62.35 and a F-score of 84.16. We found that most false negative examples are sparse entity pairs that describe a image region in an image in a very diverse way, for example “fire” and “a flaming hurdle” (Figure 4d (left)), “bananas” and “fruit” (Figure 4d (middle)). These pairs are difficult not only for using textual features, but also for using image attentions. Most false positive examples are produced by the noisy phrase embedding method. For example, “boots” and “high heels” referring to the shoes on a boy and a lady, respectively, are identified as a paraphrase pair because of their closeness in the embedding space (Figure 4d (right)). Some of the false positive examples are caused by the noise introduced by the wrong attention in an image. For example, “a green snowman” and “his new toy” are attended to the similar image regions.

7 Conclusion

In this paper, we proposed iParaphrasing: a novel task to extract VGPs describing the same visual concept in an image. We not only applied various existing techniques for this task, but also proposed a NN-based method that uses both the textual and visual information to model the similarity between the VGPs. Experiments on the Flickr30k entities dataset showed that we achieved good performance.

For future work, we plan to study a multi task method for both VGP extraction and phrase localization to further improve the performance. We worked on the Flickr30k entities dataset, where noun phrases are given and VGP supervision is available, extracting VGP in an end-to-end manner without supervision in other datasets such as the Microsoft COCO caption dataset (Chen et al., 2015) is a more realistic scenario and could be more interesting. Extracting other types of paraphrases (e.g., prepositional and verb paraphrases) is another possible extension, which requires a much deeper understanding of the relation between phrases and image regions. We also plan to apply the VGPs for CV and NLP multimodal tasks, such as VQA.

Acknowledgement

This work was supported by ACT-I, JST and JSPS KAKENHI No. 18H03264. We are very appreciated to Prof. Kumiyo Nakakoji, Prof. Yuki Arase and Prof. Sadao Kurohashi for the helpful discussion of this paper. We also thank the anonymous reviewers for their insightful comments.

References

- Ion Androutsopoulos and Prodrinos Malakasiotis. 2010. A survey of paraphrasing and textual entailment methods. *Journal of Artificial Intelligence Research*, 38(1):135–187, May.
- Yuki Arase and Jun’ichi Tsujii. 2017. Monolingual phrase alignment on parse forests. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1–11, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Colin Bannard and Chris Callison-Burch. 2005. Paraphrasing with bilingual parallel corpora. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, pages 597–604, Ann Arbor, Michigan, June. Association for Computational Linguistics.
- Regina Barzilay and Lillian Lee. 2003. Learning to paraphrase: An unsupervised approach using multiple-sequence alignment. In *Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics*, pages 16–23, Edmonton, May. Association for Computational Linguistics.
- Jonathan Berant and Percy Liang. 2014. Semantic parsing via paraphrasing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics*, pages 1415–1425, Baltimore, Maryland, June. Association for Computational Linguistics.
- Raffaella Bernardi, Ruket Cakici, Desmond Elliott, Aykut Erdem, Erkut Erdem, Nazli Ikingler-Cinbis, Frank Keller, Adrian Muscat, and Barbara Plank. 2016. Automatic description generation from images: A survey of models, datasets, and evaluation measures. *Journal of Artificial Intelligence Research*, 55(1):409–442, January.
- Rahul Bhagat and Eduard Hovy. 2013. What is a paraphrase? *Computational Linguistics*, 39(3):463–472, September.
- Rahul Bhagat and Deepak Ravichandran. 2008. Large scale acquisition of paraphrases for learning surface patterns. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: the Human Language Technology Conference*, pages 674–682, Columbus, Ohio, June. Association for Computational Linguistics.
- Peter F. Brown, Stephen A. Della Pietra, Vincent J. Della Pietra, and Robert L. Mercer. 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics*, 19(2):263–312.
- Chris Callison-Burch. 2008. Syntactic constraints on paraphrases extracted from parallel corpora. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 196–205, Honolulu, Hawaii, October. Association for Computational Linguistics.
- David Chen and William Dolan. 2011. Collecting highly parallel data for paraphrase evaluation. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 190–200, Portland, Oregon, USA, June. Association for Computational Linguistics.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C. Lawrence Zitnick. 2015. Microsoft COCO captions: Data collection and evaluation server. *CoRR*, abs/1504.00325.
- Chenhui Chu and Sadao Kurohashi. 2016. Paraphrasing out-of-vocabulary words with word embeddings and semantic lexicons for low resource statistical machine translation. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 644–648, Portoro, Slovenia, may. European Language Resources Association (ELRA).
- Volkan Cirik, Taylor Berg-Kirkpatrick, and Louis-Philippe Morency. 2018. Using syntax to ground referring expressions in natural images. In *The Thirty-Second AAAI Conference on Artificial Intelligence*.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M.F. Moura, Devi Parikh, and Dhruv Batra. 2017a. Visual Dialog. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Abhishek Das, Satwik Kottur, José M.F. Moura, Stefan Lee, and Dhruv Batra. 2017b. Learning cooperative visual dialog agents with deep reinforcement learning. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- Brendan J Frey and Delbert Dueck. 2007. Clustering by passing messages between data points. *Science*, 315(5814):972–976.

- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach. 2016. Multi-modal compact bilinear pooling for visual question answering and visual grounding. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 457–468, Austin, Texas, November. Association for Computational Linguistics.
- Ruka Funaki and Hideki Nakayama. 2015. Image-mediated learning for zero-shot cross-lingual document retrieval. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 585–590, Lisbon, Portugal, September. Association for Computational Linguistics.
- Juri Ganitkevitch and Chris Callison-Burch. 2014. The multilingual paraphrase database. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 4276–4283, Reykjavik, Iceland, May. European Language Resources Association (ELRA).
- Yunchao Gong, Qifa Ke, Michael Isard, and Svetlana Lazebnik. 2014. A multi-view embedding space for modeling internet images, tags, and their semantics. *International Journal of Computer Vision*, 106(2):210–233, Jan.
- Dan Han, Pascual Martínez-Gómez, and Koji Mineshima. 2017. Visual denotations for recognizing textual entailment. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2843–2849, Copenhagen, Denmark, September. Association for Computational Linguistics.
- David R. Hardoon, Sandor R. Szedmak, and John R. Shawe-taylor. 2004. Canonical correlation analysis: An overview with application to learning methods. *Neural Computation.*, 16(12):2639–2664, December.
- Zellig S. Harris. 1954. Distributional structure. *Word*, 10(23):146–162.
- Ronghang Hu, Marcus Rohrbach, Jacob Andreas, Trevor Darrell, and Kate Saenko. 2017. Modeling relationships in referential expressions with compositional modular networks. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July.
- Lawrence Hubert and Phipps Arabie. 1985. Comparing partitions. *Journal of Classification*, 2(1):193–218, Dec.
- Tomoyuki Kajiwara, Mamoru Komachi, and Daichi Mochihashi. 2017. Mipa: Mutual information based paraphrase acquisition via bilingual pivoting. In *Proceedings of the 8th International Joint Conference on Natural Language Processing*, pages 80–89, Taipei, Taiwan, November. Asian Federation of Natural Language Processing.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. 2014. Referitgame: Referring to objects in photographs of natural scenes. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 787–798. Association for Computational Linguistics.
- Benjamin Klein, Guy Lev, Gil Sadeh, and Lior Wolf. 2014. Fisher vectors derived from hybrid gaussian-laplacian mixture models for image annotation. *CoRR*, abs/1411.7399.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondrej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical machine translation. In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic, June. Association for Computational Linguistics.
- Chen Kong, Dahua Lin, Mohit Bansal, Raquel Urtasun, and Sanja Fidler. 2014. What are you talking about? text-to-image coreference. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June.
- Natsuda Laokulrat, Sang Phan, Noriki Nishida, Raphael Shu, Yo Ehara, Naoaki Okazaki, Yusuke Miyao, and Hideki Nakayama. 2016. Generating video description using sequence-to-sequence model with temporal attention. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 44–52, Osaka, Japan, December. The COLING 2016 Organizing Committee.
- Kenton Lee, Luheng He, Mike Lewis, and Luke Zettlemoyer. 2017. End-to-end neural coreference resolution. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 188–197, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Dekang Lin and Patrick Pantel. 2001. Dirt – discovery of inference rules from text. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '01*, pages 323–328, New York, NY, USA. ACM.

- Wang Ling, Chris Dyer, Alan W Black, and Isabel Trancoso. 2013. Paraphrasing 4 microblog normalization. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 73–84, Seattle, Washington, USA, October. Association for Computational Linguistics.
- Bill MacCartney, Michel Galley, and Christopher D. Manning. 2008. A phrase-based alignment model for natural language inference. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pages 802–811, Honolulu, Hawaii, October. Association for Computational Linguistics.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L. Yuille, and Kevin Murphy. 2016. Generation and comprehension of unambiguous object descriptions. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June.
- Yuval Marton, Chris Callison-Burch, and Philip Resnik. 2009. Improved statistical machine translation using monolingually-derived paraphrases. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 381–390, Singapore, August. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781.
- J. Mockus. 1989. *Bayesian approach to global optimization: theory and applications*. Mathematics and its applications: Soviet series. Kluwer Academic.
- Mayu Otani, Yuta Nakashima, Esa Rahtu, Heikkilä Janne, and Naokazu Yokoya. 2016. Learning joint representations of videos and sentences with web image search. In *European Conference on Computer Vision (ECCV)*, pages 651–667, October.
- Sang Phan, Yusuke Miyao, Duy-Dinh Le, and Shin’ichi Satoh. 2016. Video event detection by exploiting word dependencies from image captions. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3318–3327, Osaka, Japan, December. The COLING 2016 Organizing Committee.
- Bryan A. Plummer, Liwei Wang, Chris M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. 2015. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 2641–2649, December.
- Bryan A. Plummer, Arun Mallya, Christopher M. Cervantes, Julia Hockenmaier, and Svetlana Lazebnik. 2017. Phrase localization and visual relationship detection with comprehensive image-language cues. In *The IEEE International Conference on Computer Vision (ICCV)*, pages 1928–1937, Oct.
- Michaela Regneri, Marcus Rohrbach, Dominikus Wetzels, Stefan Thater, Bernt Schiele, and Manfred Pinkal. 2013. Grounding action descriptions in videos. *Transactions of the Association for Computational Linguistics*, 1:25–36.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, pages 91–99. Curran Associates, Inc.
- Stefan Riezler, Alexander Vasserman, Ioannis Tsochantaridis, Vibhu Mittal, and Yi Liu. 2007. Statistical machine translation for query expansion in answer retrieval. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 464–471, Prague, Czech Republic, June. Association for Computational Linguistics.
- Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. 2016. Grounding of textual phrases in images by reconstruction. In *European Conference on Computer Vision (ECCV)*, pages 817–834, October.
- Karen Simonyan and Andrew Zisserman. 2015. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*, pages 1–14.
- Matthew G. Snover, Nitin Madnani, Bonnie Dorr, and Richard Schwartz. 2009. Ter-plus: Paraphrase, semantic, and alignment enhancements to translation edit rate. *Machine Translation*, 23(2-3):117–127, September.
- Wee Meng Soon, Hwee Tou Ng, and Daniel Chung Yong Lim. 2001. A machine learning approach to coreference resolution of noun phrases. *Comput. Linguist.*, 27(4):521–544, December.
- Lucia Specia, Stella Frank, Khalil Sima’an, and Desmond Elliott. 2016. A shared task on multimodal machine translation and crosslingual image description. In *Proceedings of the First Conference on Machine Translation*, pages 543–553, Berlin, Germany, August. Association for Computational Linguistics.

- Bertrand Thirion, Edouard Duschenay, Vincent Michel, Gael Varoquaux, Olivier Grisel, Jacob VanderPlas, alexandre granfort, fabian pedregosa, Andreas Mueller, and Gilles Louppe. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3156–3164, June.
- Liwei Wang, Yin Li, and Svetlana Lazebnik. 2016a. Learning deep structure-preserving image-text embeddings. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5005–5013, June.
- Mingzhe Wang, Mahmoud Azab, Noriyuki Kojima, Rada Mihalcea, and Jia Deng. 2016b. Structured matching for phrase localization. In *European Conference on Computer Vision (ECCV)*, pages 696–711, October.
- Qi Wu, Damien Teney, Peng Wang, Chunhua Shen, Anthony Dick, and Anton van den Hengel. 2017. Visual question answering: A survey of methods and datasets. *Computer Vision and Image Understanding*, pages 1–20.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proceedings of the 32nd International Conference on Machine Learning*, volume 37 of *Proceedings of Machine Learning Research*, pages 2048–2057, Lille, France, 07–09 Jul. PMLR.
- Raymond Yeh, Jinjun Xiong, Wen-Mei W. Hwu, Minh Do, and Alexander G. Schwing. 2017. Interpretable and globally optimal prediction for textual grounding using image concepts. In *Advances in Neural Information Processing Systems*, pages 1909–1919.
- Wenpeng Yin, Hinrich Schutze, Bing Xiang, and Bowen Zhou. 2016. Abcnn: Attention-based convolutional neural network for modeling sentence pairs. *Transactions of the Association for Computational Linguistics*, 4:259–272.
- Peter Young, Alice Lai, Micah Hodosh, and Hockenmaier Julia. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association of Computational Linguistics*, 2(1):67–78.
- Liang Zhou, Chin-Yew Lin, Dragos Stefan Munteanu, and Eduard Hovy. 2006. Paraeval: Using paraphrases to evaluate summaries automatically. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, Main Conference*, pages 447–454, New York City, USA, June. Association for Computational Linguistics.