

Crisscrossed Captions: Extended Intramodal and Intermodal Semantic Similarity Judgments for MS-COCO

Zarana Parekh, Jason Baldridge, Daniel Cer, Austin Waters, and Yinfei Yang
Google Research

{zarana, jasonbaldridge, cer, austinwaters, yinfeiy}@google.com

Abstract

By supporting multi-modal retrieval training and evaluation, image captioning datasets have spurred remarkable progress on representation learning. Unfortunately, datasets have limited cross-modal associations: images are not paired with other images, captions are only paired with other captions of the same image, there are no negative associations and there are missing positive cross-modal associations. This undermines research into how inter-modality learning impacts intra-modality tasks. We address this gap with *Crisscrossed Captions* (CxC), an extension of the MS-COCO dataset with human semantic similarity judgments for **267,095** intra- and inter-modality pairs. We report baseline results on CxC for strong existing unimodal and multi-modal models. We also evaluate a multitask dual encoder trained on both image-caption and caption-caption pairs that crucially demonstrates CxC’s value for measuring the influence of intra- and inter-modality learning.

1 Introduction

Phrases such as *blue*, *chair*, and *garden path* have strong visual components, yet computational word representations are usually created with text-only corpora. Encouragingly, some recent work that derives representations using visual contexts shows improvements for both word similarity ranking and image-text retrieval (Kiros et al., 2018), and query-based training of image models demonstrates language’s power to improve image representations (Juan et al., 2020). Learning representations for both vision and language jointly should be even more effective—indeed, much progress has been made on such cross-modal learning using image captioning data (Karpathy and Li, 2015; Harwath and Glass, 2017; Faghri et al., 2018; Li et al., 2019). However, it is not yet clear whether learning representations in multimodal contexts improves perfor-

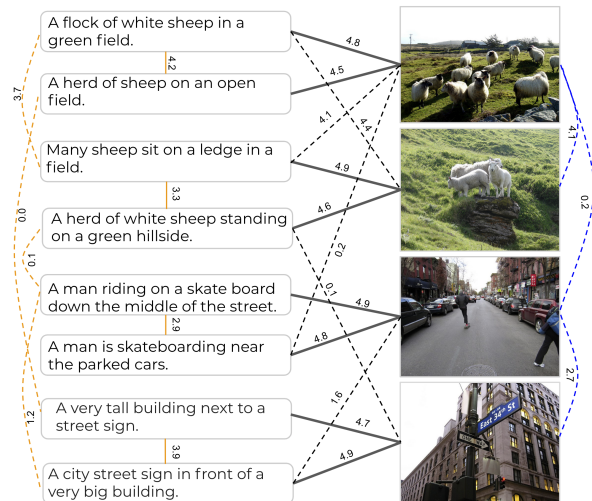


Figure 1: Crisscrossed Captions extends the MS-COCO evaluation sets by adding semantic similarity ratings for existing image-caption pairs and co-captions (solid lines), and it increases annotation density by adding further ratings for new image-caption, caption-caption and image-image pairs (dashed lines).

mance *within* as well as across modalities as there are no datasets ideally suited for this at present.

Image captioning datasets such as Flickr8k (Rashtchian et al., 2010), Flickr30k (Young et al., 2014), Multi30k (Elliott et al., 2016), Microsoft Common Objects in COntext (MS-COCO) (Lin et al., 2014), and Conceptual Captions (Sharma et al., 2018) only capture relationships between images and textual captions created for them. They miss many valid relationships between unassociated images and captions, from captions to other captions, and from images to other images. We address this gap with Crisscrossed Captions (CxC, exemplified in Figure 1), a dataset with graded, denser annotations for relationships between and among captions and images in the MS-COCO evaluation splits of (Karpathy and Li, 2015) (with 25k English captions and 5k images each).

CxC extends MS-COCO’s existing image-caption pairs with continuous (0-5) semantic similarity ratings for those pairs and new pairs. The rating criteria extend those used for Semantic Textual Similarity (Agirre et al., 2012). Intramodal pairs are selected for annotation via an indirect sampling scheme biased to gain a broad distribution of similarities. In all, CxC contains *human* ratings for **267,095** pairs (derived from 1,335,475 independent judgments), a massive extension in scale and detail to the 50k original binary pairings.

MS-COCO *incompletely* supports three retrieval tasks: image-text, text-image and text-text. CxC enhances all of these with new positive pairs, and it also supports a *new* image-image retrieval task. With its graded similarity judgements, CxC also supports correlation measures comparing model and human rankings. Retrieval metrics focus on positive pairs, but CxC’s correlation scores additionally account for low-scoring items (non matches). Supporting these evaluations on a common set of images and captions makes them more valuable for understanding inter-modal learning—compared to disjoint sets of caption-image, caption-caption, and image-image associations. Also, multimodal representations such as CLIP (Radford et al.) are useful for downstream tasks such as Visual Question Answering (Goyal et al., 2017), Vision and Language Navigation (Majumdar et al., 2020), Referring Expressions (Yu et al., 2018) and Visual Commonsense Reasoning (Zellers et al., 2019), and we hope the additional relationships and evaluations provided by CxC will help develop even better representations for tasks that span these modalities.

To establish baselines for CxC, we provide results for existing unimodal models for text (Bag-of-Words, USE (Cer et al., 2018)) and images (InceptionV3 (Szegedy et al., 2016), ResNet-152 (He et al., 2016), SimCLRv2 (Chen et al., 2020b)) as well as for two cross-modal retrieval models VSE++ (Faghri et al., 2018) and VSRN (Li et al., 2019).

We furthermore demonstrate CxC’s utility by evaluating a dual encoder that combines a bidirectional loss for image-text retrieval with a loss for text-text retrieval. The text encoder is composed of transformer layers over pre-trained BERT word representations and the image encoder is a pre-trained EfficientNet (B4) (Tan and Le, 2019a). This model delivers the strongest overall perfor-

mance across all four retrieval tasks and correlation with human scores for text-text, image-image and image-text similarity. Compared to the same dual encoder trained only with image-text pairs, this model realizes small gains for image-text tasks and large gains for text-text task but with some degradation for image-image tasks. This indicates that the model trades capacity to encode images for better text encoding—an insight that would not be easily assessed without CxC’s image-image annotations.

Our main contributions are the following:

- We describe a method for sampling items to get a broad distribution of similarities.
- We annotate the semantic similarity of 267,095 pairs. These enhance existing retrieval tasks and support a new image-image retrieval task. They also support correlation measures; these assess models’ judgments of both positive *and* negative associations.
- We establish baseline scores for existing models and a multitask dual encoder on all tasks and demonstrate that CxC allows model performance to be assessed more holistically.
- With its new positive pairs, CxC improves the recall@k measures common in image-text and text-image retrieval. This shows a 1-3% increase in recall@k over several models.
- We release CxC’s annotations at <https://github.com/google-research-datasets/Crisscrossed-Captions>, along with code to merge CxC with existing MS-COCO data.

2 Dataset Collection

Existing resources already support learning joint representations of images and text. However, we need better evaluation resources, so we extend the MS-COCO evaluation splits with graded similarity associations within and across modalities. MS-COCO has five captions for each image, split by (Karpathy and Li, 2015) into 410k training, 25k development, and 25k test captions (82k/5k/5k for images). An ideal extension would rate every pair, but this is infeasible¹ and most pairs are dissimilar anyway. To obtain new pairs with high expected similarity, we introduce a biased sampling scheme.

The data is collected in two phases. First, we define an indirect sampling scheme that uses model-based similarities from the co-modality items to

¹A split with 5k images and 25k captions has $\approx 12.5M$ image-image, $\approx 312M$ caption-caption and $\approx 125M$ image-caption pairs, so annotating all items in the validation and test splits with 5 replications would require $\approx 4.5B$ judgments.



Caption 1: A tennis player swinging a racket at a ball.

Caption 2: A man playing tennis with a crowd watching.



Caption 3: A living room with some black furniture and a colorful rug.

Caption 4: A dog laying on a leather sofa in a living room.

Figure 2: **Top:** Captions of the same image are often not paraphrases. **Bottom:** Such *co-captions* often focus on different aspects and can diverge substantially.

select *intra*modality pairs. We use these items and their human ratings to select *inter*modality pairs for annotation. We also annotate all existing *inter*modal pairs and a large sample of *co*-captions (captions associated with the same image). See the appendix for details about the annotation interface and instructions, composition of the dataset and illustrative examples.

Intramodality Two images of a man and a dog can be described differently, while two similar sentences about a man and a dog can describe dissimilar images. In Figure 2, caption 1 gives a visual description while caption 2 gives a broader event description. Divergences also occur when caption creators perceive a scene differently: caption 3 describes the room and caption 4 focuses on the dog and sofa. This semantic gap between images and their captions creates an opportunity to sample *intra*modal pairs with varying similarities. Our key idea is to use model-based similarities of images for biased sampling of caption pairs, and vice versa, and use existing image-caption pairs as pivots between modalities. This selects image pairs that are different in appearance but similar in what they depict based on their descriptions, and vice versa.

Denote the known images and captions as V ($v_1 \dots v_n$) and C ($c_1 \dots c_n$) (the latter representing *co*-caption groups of five captions each). Each item is encoded with an off-the-shelf unimodal model. Cosine similarity between items defines two symmetric matrices: S^C (pairwise caption similarities) and S^V (pairwise image similarities). The diagonals are set to zero to not sample identical items.

We encode images with Graph-RISE (486) and

construct S^I , the image-based similarity for pairs of *co*-caption groups. We encode captions with Universal Sentence Encoder (USE) (Cer et al., 2018) and average bag of words (BoW) based on GloVe embeddings (Pennington et al., 2014). *Co*-caption representations are averaged to create a single representation. From these, we construct S^C , the caption-based similarity for images pairs. USE and BoW embeddings produce two S^C matrices, but we gloss over this detail below.

We use S^C to select *image* pairs and S^I for *caption* pairs. Because of the cross-modal semantic gap, diversity and size of the underlying data, these pairs exhibit a wide range of similarity. Selecting the five most similar items (according to model-based S^V and S^C) thus produces good representation of varying amounts of similarity as judged by people. Because S^V covers *co*-caption groups, one caption is randomly chosen from each group to produce a caption pair for rating.

Caption-caption and image-image candidates are referred to as *C2C* and *I2I*, respectively. *I2I* pairs are selected with the above *other-modality* method. For *C2C* pairs, we sample half the pairs using the *other-modality* method and half from within *co*-captions. The latter introduces (mostly) positive associations between caption pairs describing the same image. This gives a balanced set of caption pairs describing same and different images.

Pairs in *C2C* and *I2I* are scored by in-house raters using a continuous scale between 0 and 5. We adopt the widely used Semantic Textual Similarity (STS) (Cer et al., 2017) for text pairs and extend it to images to define Semantic Image Similarity (SIS). To recognize that this is a graded (rather than discrete) judgment, we encouraged raters to select scores like 1.3 and obtain the final score for a pair as the average of five individual ratings.

Intermodality We select caption-image candidates *C2I* based on human ratings for *I2I* and *C2C* pairs. We mainly seek new positive matches like those identified by annotators in Ilharco et al. (2019). For each *I2I* pair (i_j, i_k) , a *C2I* pair (c_k, i_j) is generated, where c_k is a MS-COCO caption for i_k . We generate pairs from *C2C* similarly. Half of the *C2I* pairs are selected based on *C2C* ranks and the other half by *I2I* ranks (skipping pairs already selected from *C2C*). Finally, all MS-COCO pairs (25k in validation and 25k in test) are selected to obtain caption-image similarity ratings for the known items.

Caption 1	Caption 2	STS	BoW	USE
A man standing on a tennis court holding a racquet.	A man standing on a tennis court holding a tennis racquet.	5.0	0.98	0.99
A man riding a skateboard off the side of a ramp.	A man riding up the side of a skateboard ramp.	4.2	0.99	0.94
A yellow tray topped with a cup of coffee and a donut.	A white plate topped with donuts sitting on a stove top.	3.1	0.94	0.39
A bird sitting on top of a park bench.	An empty park bench sitting in front of trees.	2.2	0.90	0.69
An old car sitting on top of a lush green field.	A couple of motorcycles parked next to each other.	1.3	0.85	0.21
A man sanding next to an orange frisbee.	A couple of swans swimming in a pond next to two people.	0.2	0.84	0.11

Table 1: A comparison of CxC STS annotation scores and cosine similarity scores using GloVe BoW embeddings and Universal Sentence Encoder (USE) for five example MS-COCO caption pairs.

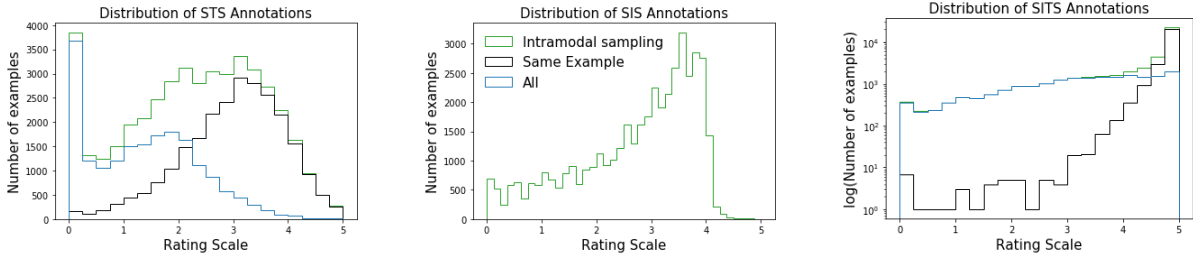


Figure 3: Distribution of ratings for the CxC validation set. *Intramodal Sampling* refers to examples selected using *other-modality* selection, *Same Example* refers to original MS-COCO pairs, and *All* covers all examples for a task.

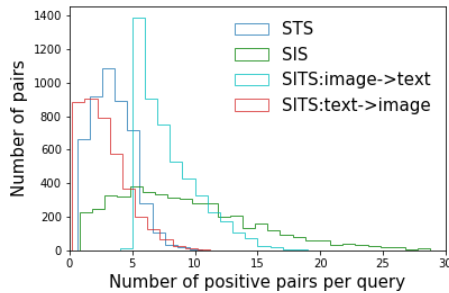


Figure 4: Distribution of counts of positive pairs (score ≥ 3) of annotations for each task (validation split).

We extend STS to define Semantic Image-Text Similarity (SITS). Raters provide a continuous score from 0 to 5 using an interface similar to that for STS and SIS. Each *C2I* pair receives five ratings; the average is used as the final SITS score.

3 Crisscrossed Captions Dataset

Using our selection and annotation methodology, we obtained ratings for 267,095 caption-caption, image-image, and caption-image pairs (1,335,475 total judgments). Figure 3 shows rating distributions for each task (validation split). It also shows the distributions of ratings for STS and SIS pairs included from other-modality selection and from original MS-COCO pairs. The test set distributions are similar. Figure 4 gives the distribution of counts of positive examples in each task (validation split), where a score ≥ 3 (for STS, SITS) and a score ≥ 2.5 (for SIS) is considered positive.

These positive examples are used for intermodal and intramodal retrieval evaluation.

STS. The majority of caption pairs selected using image similarity are negative (ratings in $[0, 3)$), which is expected given the divergences noted in Figure 2. Nevertheless, the approach produces 20,587 positive pairs. Table 1 shows pairs with their STS annotation scores and cosine similarity with BoW and USE embeddings. There is broad agreement, but the annotated similarity is not fully captured by either BoW or USE. USE provides a broader range, but scores the third pair lower than the fourth. BoW scores are bunched within a high similarity band² that aligns well with these five examples. Overall, there is a weak positive correlation between BoW and STS scores, as shown in Figure 5, which plots average BoW cosine similarity versus STS for 1000 randomly sampled pairs.

Figure 6 shows a pair of captions (and corresponding images) selected by the other-modality strategy with higher STS compared to their respective co-captions. For co-caption pairs, STS scores are more positive but many are still negative (Figure 3, left). Thus, combining both approaches leads to a more representative distribution overall. The large number of negative pairs from co-captions underscores the problem with assuming captions of the same image are paraphrases.

SIS. All image pairs *I2I* are selected using the other-modality strategy. This plus the stringent cri-

²BoW scores fall mostly in the range .8 to 1.0 over all possible pairs; STS scores fall mostly in $[.2, .8)$.

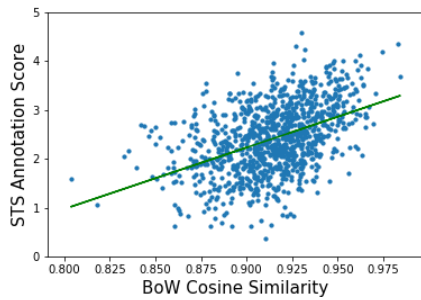


Figure 5: Plot BoW cosine similarity and STS scores for a sample of caption pairs, with the line of best fit.




Figure 6: An example of *other-modality* caption pairs (horizontal pair) with higher STS compared to their respective co-captions (vertical pairs). Each caption is placed under its corresponding image.

teria for SIS rating of 5 means very few examples are rated above 4. Nevertheless, there are many pairs with $SIS \geq 3$, indicating there are many images depicting similar scenes and events.

SITS. As shown in Figure 4, there are many more pairs with 4-5 SITS ratings, compared to STS and SIS. This is by design, as the *C2I* pairs are selected based on decreasing STS/SIS scores. This captures more positive intermodality associations and augments the existing validation and test splits. Since these pairs are missing from the existing data, they are among the examples that inappropriately penalize a model that identifies them correctly in image-caption retrieval. The SITS ratings collected for known pairs also support new correlation based evaluations, as discussed in the next section.

4 Evaluation Tasks and Metrics

CxC supports intermodal retrieval like MS-COCO but with denser annotations between image-caption pairs. It also enables intramodal retrieval and semantic similarity correlation evaluations, which were not possible before.

Karpathy and Li (2015) first used MS-COCO for image-to-caption and caption-to-image retrieval. We extend the existing associations with positive CxC pairs, and also add new caption-to-caption and image-to-image retrieval tasks using positive STS and SIS pairs (a total of four retrieval tasks). To the best of our best knowledge, CxC is the first dataset to support image-to-image retrieval over captioned images. Following Karpathy and Li (2015), we evaluate using Recall@K ($R@K$), computed as the fraction of times a correct item was found among the top K results, and median rank (med. r) of the closest ground truth result in the list.

Semantic similarity tasks such as Semantic Textual Similarity (Cer et al., 2017) and Visual Semantic Textual Similarity (vSTS) (de Lacalle et al., 2020) require a model to produce a continuous similarity score given two inputs. Typically, the models are evaluated based on the Pearson’s r of their scores with the human judgments over a set of input pairs. This is valid when training data is available to calibrate model scores to the human ratings. With CxC, we do not have such training data, so we instead use Spearman’s r to assess whether a model *rank*s pairs similarly to human raters.

It would be tempting to simply measure Spearman’s r over all pairs, but this would be flawed because CxC’s dense annotation means that the scores between many pairs are themselves correlated. To mitigate this, we use a sampled bootstrap correlation instead. For each correlation estimate, we sample half of the queries (to increase diversity across samples) and for each selected query, we choose one of the items for which CxC supplies a paired rating. We compute Spearman’s r between the CxC scores and the model scores for the selected pairs. The final correlation is the average over 1000 of these bootstrap samples.

vSTS (de Lacalle et al., 2020) contains 2677 pairs of MS-COCO captions and corresponding images. As noted above, vSTS is related dataset for multimodal semantic similarity. We considered mixing CxC and vSTS; however, this was infeasible because CxC uses the widely adopted Karpathy splits, while items in vSTS’s training, dev and test splits are spread among the Karpathy splits. We could not just make a separate cut of CxC because vSTS pairs can cross splits, e.g. an image-caption item in Karpathy training and another in Karpathy test. Given the small size of vSTS, we focused our efforts on CxC evaluations.

5 Evaluated Models

In order to establish baselines for CxC, we benchmark pretrained models for images and text. Note that these are off-the-shelf models that have not been trained on MS-COCO. We also evaluate cross-modal retrieval models that are trained on MS-COCO. Here, we focus on models that support efficient retrieval (e.g. dual encoders). We expect models with extensive cross-modal interactions, such as ViLBERT (Lu et al., 2019) and LXMERT (Tan and Bansal, 2019), will show strong performance on CxC tasks, either as standalone models that (inefficiently) score all possible item pairs or as rerankers for outputs of retrieval models.

To the best of our knowledge, there is no prior work that explores joint learning or evaluation on intra- and inter-modality retrieval tasks. Ngiam et al. (2011) and Collell Talleda and Moens (2016) show evidence that inter-modality learning helps improve intra-modality performance, but do not explore multitask learning. Lu et al. (2020) explore multitask learning but only focus on intermodal representation learning for intermodal downstream tasks. To illustrate how CxC allows us to measure how intermodal representation learning can improve both intra- and inter-modality performance, we train a dual encoder model on bidirectional image-text and text-text in-batch retrieval losses.

5.1 Pretrained Model Baselines

Text-only Models First, we use a bag-of-words (*BoW*) approach using averaged GloVe embeddings (Pennington et al., 2014) for each token in a caption as the caption representation. Second, the Universal Sentence Encoder (*USE*) (Cer et al., 2018) is a sentence level representation model that has shown strong performance on the related STS benchmark. We use the multilingual transformer version from TensorFlow Hub (Yang et al., 2020).³

Image-only Models *InceptionV3*, *ResNet-152*, and *SimCLRv2* are deep convolutional models (Szegedy et al., 2016; He et al., 2016; Chen et al., 2020a,b) trained on the ImageNet dataset. We extract 2048-dimensional image-level representations on a central crop containing 87.5% of the original image area. We access them via TensorFlow Hub.⁴

³universal-sentence-encoder-multilingual-large/1

⁴imagenet/inception_v3/feature_vector/4, imagenet/resnet_v1_152/feature_vector/4 and gs://simclr-checkpoints/simclr2/finetuned_100pct/r50_1x_sk0/hub/ respectively

Intermodal Models *VSE++* (Faghri et al., 2018) is a dual encoder (see Sec. 5.2) trained to learn a joint space of aligned images and captions. The state-of-the-art VSRN model (Li et al., 2019) is another dual encoder that uses additional training annotations to predict and use bounding boxes for more fine-grained and coherent image analysis, while using only a simple text encoder trained from scratch.⁵

5.2 Dual Encoder Baselines

We also consider several neural baseline models, all of which are *dual encoders* (Gillick et al., 2018; Yang et al., 2019) that encode both inputs separately. Dual encoder models have been proven as an effective approach to learn strong semantic representations (Cer et al., 2018; Chen et al., 2020a,b). They are often trained using an *in-batch sampled softmax* loss, as this has been observed to converge quickly and perform well on retrieval tasks (Gillick et al., 2018; Yang et al., 2019). We employ the bidirectional in-batch sampled softmax loss (eq. 1):

$$\mathcal{L} = -\frac{1}{K} \sum_{i=1}^K \left(S(l_i, r_i) - \log \sum_{j=1, j \neq i}^K e^{S(l_i, r_j)} \right) - \frac{1}{K} \sum_{i=1}^K \left(S(r_i, l_i) - \log \sum_{j=1, j \neq i}^K e^{S(r_i, l_j)} \right) \quad (1)$$

where $S(x, y)$ is the dot product of embeddings of examples x and y . This loss encourages the score of a correct pair $S(l_i, r_i)$ to be higher than scores of non-matching input pairs from the batch $S(l_i, r_j)$. Unlike full cross-attention models, this architecture enables large-scale retrieval through approximate nearest neighbor search.

We train dual encoders for caption-image and caption-caption tasks, as well as a multitask model that combines both tasks. We use EfficientNet-B4 (Tan and Le, 2019b) (pre-trained on ImageNet) as our image encoder; it yields a 1792-dimensional representation. The text encoder employs a frozen⁶ BERT-Base model (Devlin et al., 2019) followed by three transformer layers. The additional transformer layers have 8 attention heads, hidden dimension of 3072, and—like BERT-base—output 768-dimensional token-level features. We use the fea-

⁵Models available at <https://github.com/fartashfvsepp> (checkpoint "runs/coco_vse++/model_best.pth.tar") and <https://github.com/KunpengLi1994/VSRN> (checkpoint "pretrain_model/coco/model_coco.1.pth.tar").

⁶Freezing BERT makes performance slightly worse, but makes training much faster.

Annotations	Model	Image → Text				Text → Image			
		R@1	R@5	R@10	med r	R@1	R@5	R@10	med r
MS-COCO	VSE++	41.3	71.1	81.2	2	30.3	59.4	72.4	4
	VSRN-github	50.3	79.6	87.9	1	37.9	68.5	79.4	2
	VSRN-paper	53.0	81.1	89.4	-	40.5	70.6	81.1	-
	DE _{I2T}	52.0	80.3	89.5	1	37.9	67.6	78.8	2
	DE _{T2T+I2T}	54.1	81.8	89.9	1	39.7	70.0	80.9	2
CxC	VSE++	43.1	74.3	84.2	2	32.5	62.7	75.4	3
	VSRN-github	52.4	81.9	90.0	1	40.1	71.1	81.5	2
	DE _{I2T}	53.9	82.7	91.2	1	39.8	70.2	80.9	2
	DE _{T2T+I2T}	55.9	84.2	91.8	1	41.7	72.3	83.0	2

Table 2: Image ↔ Text retrieval results on the MS-COCO 5k test set and CxC’s extended pairs for the same.

Model	Text → Text			
	R@1	R@5	R@10	med r
BoW	21.2	38.2	47.4	13
USE _{mlling-large}	31.2	51.5	61.3	5
VSE++	38.7	62.3	72.2	3
VSRN-github	41.0	64.8	74.5	2
DE _{T2T}	41.7	64.4	73.4	2
DE _{I2T}	26.0	47.1	57.5	7
DE _{T2T+I2T}	42.4	64.9	74.0	2

Table 3: Text ↔ Text retrieval performance on MS-COCO 5k test set using CxC annotations.

Model	Image → Image			
	R@1	R@5	R@10	med r
InceptionV3	4.1	13.3	19.1	96
ResNet-152	11.8	35.5	49.5	11
SimCLRv2	24.5	54.9	68.1	4
VSE++	36.4	70.4	81.3	2
VSRN-github	44.2	76.7	86.2	2
DE _{I2T}	38.3	74.1	85.0	2
DE _{T2T+I2T}	38.5	73.6	84.9	2

Table 4: Image ↔ Image retrieval performance on MS-COCO 5k test set using CxC annotations.

tures at the 0^{th} token position of the final layer as the caption representation. BERT parameters are initialized from the public BERT checkpoint.⁷ The additional, trainable transformer layers are randomly initialized.

We construct three dual encoder models from these base encoders. (1) A *Text-Text* model (DE_{T2T}) uses a shared text encoder for both sides. (2) An *Image-Text* model (DE_{I2T}) uses the aforementioned text and image encoders, and includes a layer above the text encoder to project its 768 dimensions to 1792 (to match the image encoder output). (3) A *Multitask* model (DE_{T2T+I2T}) is trained on a combination of tasks (Chidambaram et al., 2019). It shares DE_{I2T}’s architecture and is trained in the same way; however, its loss is a weighted sum of image-text ($i2t$, $t2i$) and text-text ($t2t$) losses:

$$\mathcal{L} = \mathcal{L}_{i2t} + \mathcal{L}_{t2i} + c * \mathcal{L}_{t2t} \quad (2)$$

Here c is a scalar controlling the weights of losses from each task. This model has one text encoder, shared between all retrieval tasks. For hyperparameter tuning and training setup, see the appendix.

6 Results

Intermodal Retrieval Table 2 summarizes intermodal retrieval performance on both the original

MS-COCO annotations and CxC. We report performance of two versions of VSRN (Li et al., 2019)—one using the checkpoint on the author’s Github (VSRN-github, which allows us to perform CxC evaluations) and the other from the original paper (VSRN-paper, which has higher MS-COCO scores). Comparing each model on MS-COCO and CxC, the new positive items added by CxC show improved retrieval performance as they identify missing positives that are incorrectly penalized when using only original pairs (as noted in Ilharco et al. (2019) for Flickr8k). Multitask training in DE_{T2T+I2T} provides a boost over using only intermodal pairs for training (i.e. DE_{I2T}). It performs similarly with VSRN-paper—it seems likely that VSRN’s greater investment on the image analysis (with representations based on extracted object bounding boxes) is matched by DE_{T2T+I2T}’s greater investment in the text encoder.

Figure 7 shows three examples of images retrieved for caption queries. The CxC annotations capture missing examples in the first two cases, and the last two show there are still more positive pairs that remain unassociated in CxC. Figure 8 shows the same for captions retrieved from image queries, again showing that many examples are captured in CxC that are missing in MS-COCO.

Intramodal Retrieval Tables 3 and 4 give intramodal retrieval results enabled by CxC’s STS

⁷bert_en_uncased_L-12_H-768_A-12/2

Model	STS	SIS	SITS
	avg \pm std	avg \pm std	avg \pm std
BoW	55.1 \pm 0.6	–	–
USE _{mlling-large}	71.4 \pm 0.4	–	–
Inception V3	–	19.6 \pm 1.9	–
ResNet-152	–	59.2 \pm 1.3	–
SimCLRv2	–	74.3 \pm 0.9	–
VSE++	74.4 \pm 0.4	73.3 \pm 0.9	55.2 \pm 1.5
VSRN-github	73.0 \pm 0.4	70.1 \pm 1.0	60.4 \pm 1.3
DE _{T2T}	72.9 \pm 0.4	–	–
DE _{I2T}	50.9 \pm 0.6	81.3 \pm 0.7	61.6 \pm 1.4
DE _{T2T+I2T}	74.2 \pm 0.4	74.5 \pm 0.9	61.9 \pm 1.3

Table 5: Spearman’s R Bootstrap Correlation ($\times 100$) on MS-COCO 5k test set **using CxC annotations**.

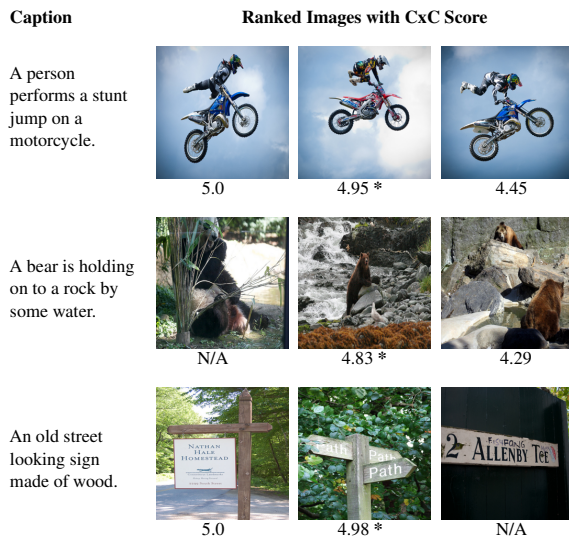


Figure 7: Text \rightarrow Image retrieval examples (MS-COCO val set), with CxC SITS score provided when available. Images are ranked from left to right, based on DE_{T2T+I2T} scores. MS-COCO annotations only consider images marked by * to be correct retrieval results.

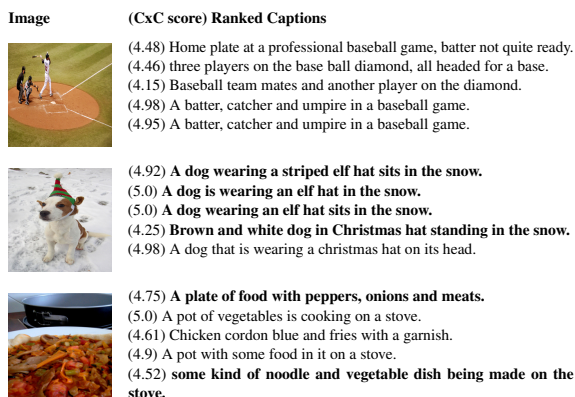


Figure 8: Image \rightarrow Text retrieval results (MS-COCO val set), ranked top to bottom using DE_{T2T+I2T} scores. MS-COCO annotations only consider the **bold** captions to be correct results.

and SIS ratings respectively. USE_{mlling-large} is a strong baseline for Text \rightarrow Text, but all the cross-modal models beat USE_{mlling-large} by a wide margin, likely due to learning on in-domain captions. InceptionV3 and ResNet-152 prove surprisingly weak for Image \rightarrow Image, but SimCLRv2 proves to be a strong unimodal baseline for this task. The cross-modal models nevertheless beat SimCLRv2 by a wide margin, even though none were trained on image-image retrieval directly. In terms of joint intra- and inter-modal learning, the multitask DE_{T2T+I2T} model provides strong, balanced performance: it is close to DE_{T2T} for Text \rightarrow Text and DE_{I2T} for Image \rightarrow Image and far outperforms the latter for Text \rightarrow Text. The strong performance is especially notable considering that both DE_{I2T} and DE_{T2T+I2T} have the same model capacity.

Semantic Similarity Table 5 shows Spearman’s R bootstrapped correlation for all models with respect to CxC’s STS, SIS and SITS scores. Overall, VSE++, VSRN-github and DE_{T2T+I2T} perform better than unimodal baselines, but interesting further patterns emerge. Despite being much worse for retrieval, VSE++ actually beats VSRN-github on STS and SIS; however, its low SITS score indicates it fails to bridge the two modalities as well. The correlation scores also show that DE_{I2T} is too focused on images: it has the highest SIS (81.3), but has worse STS (50.9) than even BoW (55.1). Adding the text-text loss to DE_{I2T} training, i.e. DE_{T2T+I2T}, produces much more balanced overall performance. On SIS, SimCLRv2 is stronger than all cross-modal models, except DE_{I2T}. SITS scores appear to rank all models similarly to retrieval (Table 2).

The fact that DE_{T2T+I2T} is better than both DE_{T2T} and unimodal baselines for STS and Text \rightarrow Text retrieval is encouraging, and it demonstrates the value of having a single set of annotations covering the relatedness of a common set of images and captions. We expect that a multitask model which also uses image-image training pairs could demonstrate gains across all tasks—measurements made possible by the CxC annotations (especially the new image-image associations).

7 Conclusion

The CxC dataset provides a much more complete set of relationships between and among images and captions than the raw MS-COCO image-caption pairs. We demonstrate that a dual encoder that learns from both image-caption pairs and caption-

caption pairs (DE_{T_{2T}+I_{2T}}) exhibits strong, balanced performance across four retrieval tasks and three correlation measures. There is much remaining headroom for future models for all these tasks.

CxC’s annotations themselves validate the strong semantic alignment between images and their original captions—these have an average similarity of 4.85. However, we also find that co-captions (captions for the same image) have an average score of just 3.0. This calls into question the use of such pairs in training and evaluating paraphrase generation models (Gupta et al., 2018) and reinforces the need for images as context for human evaluation in paraphrasing (Wang et al., 2019).

Acknowledgments

We thank Julia Hockenmaier for her inputs on CxC’s formulation, the Google Data Compute Team, especially Ashwin Kakarla and Mohd Ma-jeed for their tooling and annotation support, Yuan Zhang, Eugene Ie for their comments on the initial versions of the paper and Daphne Luong for executive support for the data collection.

References

- Eneko Agirre, Daniel Cer, Mona Diab, and Aitor Gonzalez-Agirre. 2012. Semeval-2012 task 6: A pilot on semantic textual similarity. In ** SEM 2012: The First Joint Conference on Lexical and Computational Semantics—Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 385–393.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St. John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, Brian Strope, and Ray Kurzweil. 2018. [Universal sentence encoder for English](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 169–174, Brussels, Belgium. Association for Computational Linguistics.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey E. Hinton. 2020a. [A simple frame-work for contrastive learning of visual representations](#). *CoRR*, abs/2002.05709.
- Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey E. Hinton. 2020b. [Big self-supervised models are strong semi-supervised learners](#). *CoRR*, abs/2006.10029.
- Muthu Chidambaram, Yinfei Yang, Daniel Cer, Steve Yuan, Yunhsuan Sung, Brian Strope, and Ray Kurzweil. 2019. [Learning cross-lingual sentence representations via a multi-task dual-encoder model](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepLanLP-2019)*, pages 250–259, Florence, Italy. Association for Computational Linguistics.
- Guillem Collell Talleda and Marie-Francine Moens. 2016. Is an image worth more than a thousand words? on the fine-grain semantic differences between visual and linguistic representations. In *Proceedings of the 26th International Conference on Computational Linguistics*, pages 2807–2817. ACL.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Desmond Elliott, Stella Frank, Khalil Sima’an, and Lucia Specia. 2016. [Multi30K: Multilingual English-German image descriptions](#). In *Proceedings of the 5th Workshop on Vision and Language*, pages 70–74, Berlin, Germany. Association for Computational Linguistics.
- Fartash Faghri, David J. Fleet, Jamie Ryan Kiros, and Sanja Fidler. 2018. [VSE++: improving visual-semantic embeddings with hard negatives](#). In *British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018*, page 12. BMVA Press.
- Daniel Gillick, Alessandro Presta, and Gaurav Singh Tomar. 2018. End-to-end retrieval in continuous space. *arXiv preprint arXiv:1811.08008*.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Ankush Gupta, Arvind Agarwal, Prawaan Singh, and Piyush Rai. 2018. A deep generative framework for paraphrase generation. In *Proceedings of AAAI-18*.
- David Harwath and James Glass. 2017. [Learning word-like units from joint audio-visual analysis](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long*

- Papers*), pages 506–517, Vancouver, Canada. Association for Computational Linguistics.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Gabriel Ilharco, Yuan Zhang, and Jason Baldridge. 2019. Large-scale representation learning from visually grounded untranscribed speech. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- Da-Cheng Juan, Chun-Ta Lu, Zhen Li, Futang Peng, Aleksei Timofeev, Yi-Ting Chen, Yaxi Gao, Tom Duerig, Andrew Tomkins, and Sujith Ravi. 2020. Ultra fine-grained image semantic embedding. In *Proceedings of the 13th International Conference on Web Search and Data Mining, WSDM '20*, page 277–285, New York, NY, USA. Association for Computing Machinery.
- Andrej Karpathy and Fei-Fei Li. 2015. Deep visual-semantic alignments for generating image descriptions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3128–3137.
- Jamie Kiros, William Chan, and Geoffrey Hinton. 2018. Illustrative language understanding: Large-scale visual grounding with image search. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 922–933.
- Oier Lopez de Lacalle, Ander Salaberria, Aitor Soroa, Gorka Azkune, and Eneko Agirre. 2020. Evaluating multimodal representations on visual semantic textual similarity. In *Proceedings of the Twenty-third European Conference on Artificial Intelligence (ECAI)*, Santiago Compostela, Spain.
- Kunpeng Li, Yulun Zhang, Kai Li, Yuanyuan Li, and Yun Fu. 2019. Visual semantic reasoning for image-text matching. In *The IEEE International Conference on Computer Vision (ICCV)*.
- 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 H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 13–23. Curran Associates, Inc.
- Jiasen Lu, Vedanuj Goswami, Marcus Rohrbach, Devi Parikh, and Stefan Lee. 2020. 12-in-1: Multi-task vision and language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10437–10446.
- Arjun Majumdar, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, and Dhruv Batra. 2020. Improving vision-and-language navigation with image-text pairs from the web.
- Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. 2011. Multimodal deep learning. In *ICML*.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. *Image*, 2:T2.
- Cyrus Rashtchian, Peter Young, Micah Hodosh, and Julia Hockenmaier. 2010. Collecting image annotations using amazon’s mechanical turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 139–147. Association for Computational Linguistics.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of ACL*.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826.
- Hao Tan and Mohit Bansal. 2019. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5100–5111, Hong Kong, China. Association for Computational Linguistics.
- Mingxing Tan and Quoc Le. 2019a. EfficientNet: Rethinking model scaling for convolutional neural networks. volume 97 of *Proceedings of Machine Learning Research*, pages 6105–6114, Long Beach, California, USA. PMLR.
- Mingxing Tan and Quoc Le. 2019b. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114.

- Su Wang, Rahul Gupta, Nancy Chang, and Jason Baldridge. 2019. [A task in a suit and a tie: Paraphrase generation with semantic augmentation](#). In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pages 7176–7183. AAAI Press.
- Yinfei Yang, Daniel Cer, Amin Ahmad, Mandy Guo, Jax Law, Noah Constant, Gustavo Hernandez Abrego, Steve Yuan, Chris Tar, Yun-hsuan Sung, Brian Strope, and Ray Kurzweil. 2020. [Multilingual universal sentence encoder for semantic retrieval](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 87–94, Online. Association for Computational Linguistics.
- Yinfei Yang, Gustavo Hernandez Abrego, Steve Yuan, Mandy Guo, Qinlan Shen, Daniel Cer, Yun-hsuan Sung, Brian Strope, and Ray Kurzweil. 2019. [Improving multilingual sentence embedding using bi-directional dual encoder with additive margin softmax](#). In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*, pages 5370–5378. International Joint Conferences on Artificial Intelligence Organization.
- 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.
- Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. 2018. [Mattnet: Modular attention network for referring expression comprehension](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1307–1315.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6720–6731.

CxC Annotations - Collection and Analysis

We present additional details of the human annotation process. To annotate the CxC dataset, in house annotators were employed: 170 (74 men, 96 women) for STS, 61 (28 men, 33 women) for SIS and 113 (46 men, 67 women) for SITS. All were aged between 20-35 years. The annotators were paid hourly wages that are competitive for their locale. They have standard rights as contractors. They were fluent English speakers.

We define separate annotation interfaces for each of Semantic Textual Similarity (STS), Semantic Image Similarity (SIS) and Semantic Image Text Similarity (SITS) tasks. We define a similarity scale ranging from 0 to 5 for all three tasks, following Cer et al. (2017).

We conducted a few pilot annotation rounds with the annotators to evaluate the effectiveness of the annotation instructions and the annotation interface. We learned that allowing the annotators to rate on a continuous 0-5 scale instead of a discrete one like STS resulted in higher correlation between the individual ratings. As a result, we decided to use the continuous ratings in the task but still keep the similarity definition for each discrete value in the annotation instructions. The final annotation interfaces are illustrated in Figures 4 for STS, 5 for SIS and 6 for SITS. Task-specific high-level instructions are displayed at the top in an expandable text box followed by a pair of examples. At the bottom there is a sliding bar with 0-5 score instructions along the scale. Since the SITS instructions are longer, they are shown when the annotator hovers over the corresponding score to improve readability for this task.

Each annotator is required to evaluate the displayed example based on the instructions and score them. The sliding scale makes it intuitive for the annotators to rate an example 2.87 if they feel the semantic similarity of the pair lies between score descriptions of 2 and 3, leaning towards 3. Finally, the annotator response is recorded when they click the submit button at the bottom of the page. The absolute score is deliberately not displayed so as not to distract the workers towards trying to get a clean integer value like 3.0 instead of 2.94 or 2.97.

The annotators were able to get a better grasp of the task through the pilot annotations and got quicker at scoring the pairs. They took an average of 37, 17 and 17 seconds per example for STS, SIS and SITS tasks respectively for the final round

Task	STS	SIS	SITS	Total (per split)
Validation	44,009	42,767	44,722	131,498
Test	44,045	46,719	44,833	135,597
Total (per task)	88,054	89,486	89,555	267,095

Table 1: Number of annotations per task and split.




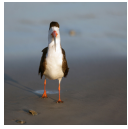








	Image Pairs - SIS annotations	Image Text Pairs - SITS annotations.
5		 A man poses with a surfboard on a beach.
4		 A couple of birds that are walking on some sand.
3		 A man is riding a surfboard at the beach.
2		 Three people stand on an empty beach watching a bird in the sky.
1		 A man in a hat rides an elephant in a river.
0		 Long road with a sign titled Jackson River Rd and East Main St.

Figure 1: Examples for each annotation score (0-5) of SIS (left) and SITS (right) tasks.

of annotations. Table 2 describes the instructions shared with the annotation workers for each task. The side-by-side comparison shows how each rating on the SIS and SITS scales compares to the STS benchmark. Figure 1 shows a set of SIS and SITS examples for the 0-5 rating scale shared along with the instructions. Table 1 contains the breakdown of the number of annotations per task per split.

Figure 2 shows a distribution of the standard deviation of raw annotations for each item per task. For STS, there is larger overall deviation compared to the other two tasks—it seems that pairs of short captions leave more ambiguity and are open for broader interpretation than when at least one image is involved. Note also that SITS is expected to have lower deviation because of the sampling based on STS and SIS annotations.

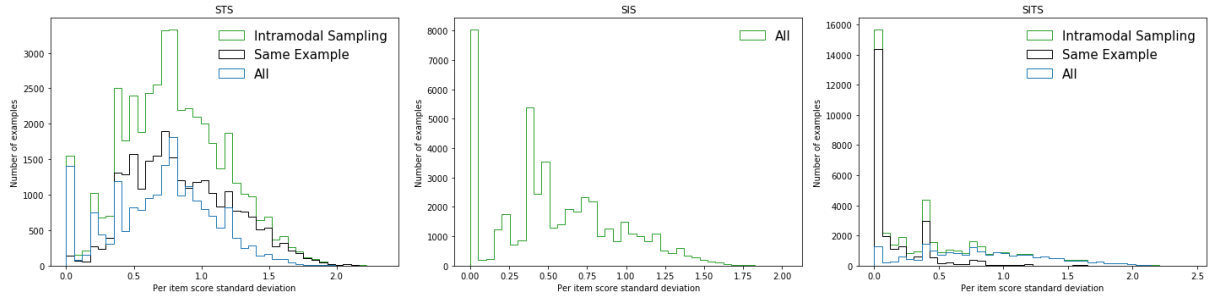


Figure 2: Standard deviation per item in CxC. *Intramodal Sampling* refers to examples selected using *other-modality* selection, *Same Example* refers to original MS-COCO pairs, and *All* covers all examples for a task.

Training Setup

Figure 3 shows the basic architecture of the dual encoder models from Section 5.2, which establish strong baselines on all the retrieval and correlation tasks. The image encoder and text encoder are both pre-trained in all experiments. Following Iharco et al. (2019), we pretrain our dual encoders on the Conceptual Captions dataset (Sharma et al., 2018) with image-to-caption and caption-to-image losses. Conceptual Captions contains 3.3 million pairs of images and captions—far larger than MS-COCO. Pre-training uses the Adam optimizer ($\beta_1 = 0.9, \beta_2 = 0.999$) and a learning rate that starts at $1e-4$ and decays by 0.1% every 1000 steps. We stop pre-training after $\approx 30k$ steps and select the checkpoint that maximizes R@10 on a held-out set. We then fine-tune this checkpoint on MS-COCO using the same hyper parameters, except for a smaller learning rate of $5e-6$.

Our models are trained on 32-core slices of Cloud TPU V3 pods, with a per-replica batch size of $K = 64$ during both pre-training and fine-tuning. Because in-batch sampled softmax loss is known to perform best when computed over a large number of negative samples (Gillick et al., 2018), our training setup pools image and caption encodings from all replicas before computing the loss. That is, each replica computes l and r for its local mini-batch and broadcasts them to all others to be used as negative samples. Training with N cores thus allows the loss to be computed over the global batch of $N \cdot K$ examples and $(N \cdot K)^2$ pairs (in our case 2048 examples and 2048^2 example pairs).

Ablation Experiments

Our model architecture and training setup differ from prior work in key ways. In particular, best known results for VSE++ and VSRN are from models that were trained with much smaller batch sizes,

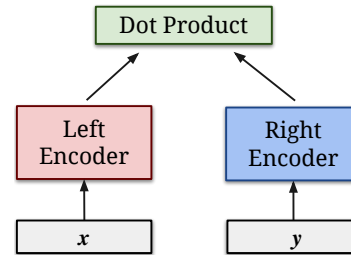


Figure 3: Dual Encoder with inputs x and y , encoded by the left and right encoders, respectively. Similarity is computed as the dot product of the encodings.

did not undergo Conceptual Captions pre-training, and had different image encoder architectures. To evaluate the effect of these factors, we trained variants of our DE_{I2T} model (here, the *baseline* training recipe) with the following one-off ablations:

- The *small batch size* ablation reduces the training batch size to 128 examples, to match that of VSE++ and VSRN in (Faghri et al., 2018) and (Li et al., 2019), respectively.
- *No pretraining* skips dual encoder pretraining on Conceptual Captions.
- *ResNet-152* uses the same recipe as *baseline*, but replaces the EfficientNet-B4 image encoder with ResNet-152, which was used in VSE++. Notably, EfficientNet-B4 has fewer parameters than ResNet-152, but achieves higher classification accuracy on ImageNet.

Table 3 summarizes the performance of the ablated models. Reducing the batch size causes a small but consistent reduction in recalls across all tasks. Removing Conceptual Captions pretraining leads to larger regressions on all tasks – except on Text-Text retrieval, where results are curiously better than the baseline. Likewise, models using ResNet-152 image encoders perform worst overall,

Score	Semantic Textual Similarity (STS)	Semantic Image Similarity (SIS)	Semantic Image-Text Similarity (SITS)
5	The texts are completely equivalent as they mean the same thing.	The scenes are near duplicates , possibly being viewed from a different perspective.	The image and sentence are perfectly matched. The sentence is an almost perfect description for the image.
4	The texts are mostly equivalent but some unimportant details differ.	The two scenes are mostly equivalent , but some unimportant details differ such as involving different but the same or highly similar types of participants, actions, objects and background.	The image and sentence are mostly matched , but some unimportant details differ such as involving different but the same or highly similar types of participants, actions, objects and background. The text can partially describe the image.
3	The texts are roughly equivalent but some important information differs or is missing.	The two scenes are roughly equivalent , but some important details are different or missing such as involving a notable difference in the types of participants, actions, objects or background.	The image and sentence are roughly matched , but some important details are different or missing such as involving a notable difference in the types of participants, actions, objects or background. The image cannot be described using the text.
2	The texts are not equivalent but share some details.	The two scenes are not equivalent , but share some details in terms of the types of participants, actions, objects or background.	The image and sentence are not matched , but share some details in one or more of the types of participants, actions, objects or background.
1	The texts are not equivalent but are on the same topic.	The two scenes are not equivalent , but are loosely thematically related.	The image and sentence are not matched , but are loosely thematically related.
0	The texts are on different topics.	The two scenes are completely dissimilar.	The image and sentence are completely unmatched.

Table 2: Intramodality annotation criteria for Semantic Image Similarity (SIS) and Intermodality annotation criteria for Semantic Image-Text Similarity (SITS) with comparison to equivalent Semantic Textual Similarity (STS) annotations (Agirre et al., 2012).

but also perform (slightly) better than the baseline on Text-Text retrieval.

Overall, we conclude that pretraining and choice of image encoder architecture have large effects on model performance; large-batch training is beneficial, but has a smaller impact. Finally, the asymmetric shifts in task performance suggest models make implicit trade-offs based on the relative difficulty of each task – here, apparently, a function of encoder strength and quantity of training data. Understanding these dynamics, and building models that perform well across all tasks, requires future study. Crisscrossed Captions enables such work by giving a more complete picture of model quality on both intra- and inter-modal tasks.

Annotations	Model	Image → Text		Text → Image		Text → Text		Image → Image	
		R@1	R@10	R@1	R@10	R@1	R@10	R@1	R@10
MS-COCO	DE _{I2T} (baseline)	52.0	89.5	37.9	78.8	25.9	57.2	–	–
	DE _{I2T} (small batch size)	49.6	88.2	35.7	78.0	25.3	57.6	–	–
	DE _{I2T} (no pretraining)	45.0	86.0	31.2	74.7	34.5	67.2	–	–
	DE _{I2T} (ResNet-152)	43.5	83.0	28.9	71.4	28.2	60.2	–	–
CxC	DE _{I2T} (baseline)	53.9	91.2	39.8	80.9	26.0	57.5	38.3	85.0
	DE _{I2T} (small batch size)	51.8	90.1	37.7	80.3	25.4	57.9	38.0	84.3
	DE _{I2T} (no pretraining)	47.0	88.1	33.2	77.6	34.6	67.6	37.0	84.0
	DE _{I2T} (ResNet-152)	45.2	85.1	30.8	74.4	28.3	60.5	29.7	76.5

Table 3: Ablation analysis for DE_{I2T}, retrieval results on MS-COCO 5k test set and CxC. (Note that MS-COCO does not support Image → Image retrieval evaluation at all.)

Instruction

For this task, you will compare two sentence length snippets of English text and rate how similar they are in meaning to each other

– Detailed Instructions (click to show/hide)

Task Instructions

Two snippets of text can mean the same thing even if they use very different words and phrases. Conversely, two texts that are superficially very similar in their word choice, phrasing and overall composition can have very different meanings.

For this task, you will compare two phrase or sentence length segments of text and select whether or not they have same underlying meaning or message in terms of what they refer to, say or ask about the world.

For example, do both snippets refer to the exact same person, action, event, idea or thing? Or are they similar but different according to either large or small details?

Tips

- Assign ratings according to the underlying meaning of the two snippets rather than their superficial similarities or differences.
- Be careful of wording differences that have an important impact on what is being said or described.
- Ignore grammatical errors and awkward wordings as long as they do not obscure what is being conveyed.
- Use the full range of scores and do not be concerned about selecting precise values such as (2.0) or (3.5).
- Be fast! Do not spend time obsessing over one numerical value versus another (e.g. a (1) vs a (2)) -- in such cases, pick a value somewhere between them and closer to the value you think represents their similarity.

Instruction

For this task, you will compare two sentence length snippets of English text and rate how similar they are in meaning to each other

+ Detailed Instructions (click to show/hide)

Task

Are the following two snippets of text equivalent in meaning?

sentence 1: A dog riding on a skateboard next to a person.

sentence 2: a man riding a skateboard with his dog.

0- On different topics

1- Not equivalent, but are on the same topic

2- Not equivalent, but share some details

3- Roughly equivalent- but some important information differs or is missing

4- Mostly equivalent- but some unimportant details differ

5- Completely equivalent as they mean the same thing

Comments (Optional)

Submit

Figure 4: STS Annotation Interface

Task: Evaluate the 2 images and answer the question below.
 More instructions on how to complete the task are available in this [guidelines doc](#).

— Detailed Instructions (click to show/hide)

For this task, you will compare two images and select whether or not they have the same underlying meaning in terms of what they refer to, say or ask about the world.

For example, do both images refer to the exact same person, action, event, idea or thing? Or are they similar but different according to either large or small visual details?

Tips

1. Assign ratings according to the underlying meaning and/or theme of the images and not only based on visual similarities or differences.
2. Use the full range of scores and do not be concerned about selecting precise values such as (2.0) or (3.5).
3. Be fast! Do not spend time obsessing over one numerical value versus another (e.g. a (1) vs a (2)) – in such cases, pick a value somewhere between them and closer to the value you think represents their similarity.

Image 1






Image 2



Do you think the 2 images above are the same?



(0) The two scenes are completely dissimilar.	(1) The two scenes are not equivalent,	(2) The two scenes are partially similar,	(3) The two scenes are roughly equivalent,	(4) The two scenes are mostly equivalent,	(5) The scenes are completely equivalent,
---	--	---	--	---	---

[Submit](#)

Figure 5: SIS Annotation Interface

Task : Compare the image and snippet of text to answer the question below
 More instruction on how to complete the task are available in this [guidelines doc](#).


— Detailed Instructions (click to show/hide)

For this task, you will compare an image and a snippet of text and select whether or not they have the same underlying meaning in terms of what they refer to, say or ask about the world.

For example, do both the image and the text snippet refer to the exact same person, action, event, idea or thing? Or are they similar but different according to either large or small details?

Tips


1. Assign ratings according to the underlying meaning and/or theme of the image and the text snippet and not only based on superficial similarities or differences.
2. Be careful of wording differences that have an important impact on what is being said or described.
3. Ignore grammatical errors and awkward wordings as long as they do not obscure what is being conveyed.
4. Use the full range of scores and do not be concerned about selecting precise values such as (2.0) or (3.5).
5. Be fast! Do not spend time obsessing over one numerical value versus another (e.g. a (1) vs a (2)) – in such cases, pick a value somewhere between them and closer to the value you think represents their similarity.



snippet of text

a couple of men standing next to each other

Do you think the image and text snippet are the same?



(0) Completely unmatched	(1) Thematically matched	(2) The image and sentence are not matched, but share some details in one or more of the types of participants, actions, objects or background.	(3) Roughly matched	(4) Mostly matched	(5) Perfectly matched
--------------------------	--------------------------	---	---------------------	--------------------	-----------------------

[Submit](#)

Figure 6: SITS Annotation Interface