

# Evaluating Multimodal Representations on Sentence Similarity: vSTS, Visual Semantic Textual Similarity Dataset

Oier Lopez de Lacalle

Aitor Soroa

Eneko Agirre

IXA NLP Group

UPV/EHU University of the Basque Country

Donostia, Basque Country

{oier.lopezdelacalle,a.soroa,e.agirre}@ehu.eus

## 1. Introduction

The success of word representations (embeddings) learned from text has motivated analogous methods to learn representations of longer sequences of text such as sentences, a fundamental step on any task requiring some level of text understanding [13]. Sentence representation is a challenging task that has to consider aspects such as compositionality, phrase similarity, negation, etc. In order to evaluate sentence representations, intermediate tasks such as Semantic Textual Similarity (STS) [2] or Natural Language Inference (NLI) [1] have been proposed, with STS being popular among unsupervised approaches<sup>1</sup>. Through a set of campaigns, STS has produced several manually annotated datasets, where annotators measure the similarity among sentences, with higher scores for more similar sentences, ranging between 0 (no similarity) to 5 (semantic equivalence). Human annotators exhibit high inter-tagger correlation in this task.

In another strand of related work, tasks that combine representations of multiple modalities have gained increasing attention, including image-caption retrieval, video and text alignment, caption generation, and visual question answering. A common approach is to learn image and text embeddings that share the same space so that sentence vectors are close to the representation of the images they describe [3, 7]. [9] provides an approach that learns to align images with descriptions. Joint spaces are typically learned combining various types of deep learning networks such as recurrent networks or convolutional networks, with some attention mechanism [19, 11, 15].

The complementarity of visual and text representations for improved language understanding have been shown also on word representations, where embeddings have been combined with visual or perceptual input to produce grounded representations of words [4, 6, 10, 8, 12, 18, 20].

These improved representation models have outperformed traditional text-only distributional models on a series of word similarity tasks, showing that visual information coming from images is complementary to textual information.

In this paper we present Visual Semantic Textual Similarity (vSTS), a dataset which allows to study whether better sentence representations can be built when having access to corresponding images, e.g. a caption and its image, in contrast with having access to the text alone. This dataset is based on a subset of the STS benchmark [2], more specifically, the so called STS-images subset, which contains pairs of captions. Note that the annotations are based on the textual information alone. vSTS extends the existing subset with images, and aims at being a standard dataset to test the contribution of visual information when evaluating sentence representations.

In addition we show that the dataset allows to explore two hypothesis: H1) whether the image representations alone are able to predict caption similarity; H2) whether a combination of image and text representations allow to improve the text-only results on this similarity task.

## 2. The vSTS dataset

The dataset is derived from a subset of the caption pairs already annotated in the Semantic Textual Similarity Task (see below). We selected some caption pairs with their similarity annotations, and added the images corresponding to each caption. While the human annotators had access to only the text, we provide the system with both the caption and corresponding image, to check whether the visual representations can be exploited by the system to solve a text understanding and inference task.

As the original dataset contained captions referring to the same image, and the task would be trivial for pairs of the same image, we filtered those out, that is, we only consider caption pairs that refer to different images. In total, the dataset comprises 829 instances, each instance containing a

<sup>1</sup>See for instance the models evaluated on STS benchmark <http://ixa2.si.ehu.es/stswiki/index.php/STSbenchmark>

subset	#pairs	mean sim	std sim	#zeroes
2014	374	1.77	1.49	78
2015	445	1.69	1.44	81
Total	819	1.72	1.46	159

Table 1. Main statistics of the dataset.




Figure 1. Similarity distribution of the visual STS dataset.

pair of images and their description, as well as a similarity value that ranges from 0 to 5. The instances are derived from the following datasets:

**Subset 2014** This subset is derived from the Image Descriptions dataset which is a subset of the PASCAL VOC-2008 dataset [17]. PASCAL VOC-2008 dataset consists of 1,000 images and has been used by a number of image description systems. In total, we obtained 374 pairs (out of 750 in the original file).

**Subset 2015** The subset is derived from Image Descriptions dataset, which is a subset of 8k-picture of Flickr. 8k-Flickr is a benchmark collection for sentence-based image description, consisting of 8,000 images that are each paired with five different captions which provide clear descriptions of the salient entities and events. We obtained 445 pairs (out of 750 in the original).

**Score distribution** Due to the caption pairs are generated from different images, strong bias towards low scores is expected (see Figure 1). We measured the score distribution in the two subsets separately and jointly, and see that the two subsets follow same distribution. As expected, the most frequent score is 0 (Table 1), but the dataset still shows wide range of similarity values, with enough variability.

### 3. Experiments

**Experimental setting** We split the vSTS dataset into development and test partitions, sampling 50% at random, while preserving the overall score distributions. In addition, we used part of the text-only STS benchmark dataset as a training set, discarding the examples that overlap with vSTS.

**STS Models** We checked four models of different complexity and modalities. The baseline is a word overlap model (OVERLAP), in which input texts are tokenized with white space, vectorized according to a word index, and similarity is computed as the cosine of the vectors. We also calculated the centroid of Glove word embeddings [16] (CAVERAGE) and then computed the cosine as a second text-based model.

Model	10-fold xval on dev set			Test set		
A - OVERLAP	0.68			0.64		
B - CAVERAGE	0.65			0.67		
C - DAM	0.71			0.69		
D - RESNET50	0.63			0.61		
Combination	LR	$\oplus$	$\otimes$	LR	$\oplus$	$\otimes$
A+D	0.77	0.77	0.77	0.76	0.75	0.75
B+D	0.75	0.73	0.70	0.76	0.73	0.70
C+D	0.78	0.78	0.78	0.77	0.77	0.78

Table 2. Pearson correlation  $r$  results in development and test. Note that A, B and C are text-only, and D is image-only.

The third text-based model is the state of the art Decomposable Attention Model [14] (DAM), trained on the STS benchmark dataset as explained above. Finally, we use the top layer of a pretrained resnet50 model [5] to represent the images associated to text, and use the cosine for computing the similarity of a pair of images (RESNET50).

**Model combinations** We combined the predictions of text based models with the predictions of the image based model (see Table 2 for specific combinations). Models are combined using addition ( $\oplus$ ), multiplication ( $\otimes$ ) and linear regression (LR) of the two outputs. We use 10-fold cross-validation on the development test for estimating the parameters of the linear regressor.

**Results** Table 2 shows the results of the single and combined models. Among single models, as expected, DAM obtains the highest Pearson correlation ( $r$ ). Interestingly, the results show that images alone are valid to predict caption similarity (0.61  $r$ ). Results also show that image and sentence representations are complementary, with the best results for a combination of DAM and RESNET50 representations. These results confirm our hypotheses, and more generally, show indications that in systems that work with text describing the real world, the representation of the real world helps to better understand the text and do better inferences.

### 4. Conclusions and further work

We introduced the vSTS dataset, which contains caption pairs with human similarity annotations, where the systems can also access the actual images. The dataset aims at being a standard dataset to test the contribution of visual information when evaluating the similarity of sentences.

Experiments confirmed our hypotheses: image representations are useful for caption similarity and they are complementary to textual representations, as results improve significantly when two modalities are combined together.

In the future we plan to re-annotate the dataset with scores which are based on both the text and the image, in order to shed light on the interplay of images and text when understanding text.

## Acknowledgments

This research was partially supported by the Spanish MINECO (TUNER TIN2015-65308-C5-1-R and MUSTER PCIN-2015-226).

## References

- [1] S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326*, 2015. [1](#)
- [2] D. Cer, M. Diab, E. Agirre, I. Lopez-Gazpio, and L. Specia. 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, August 2017. Association for Computational Linguistics. [1](#)
- [3] F. Faghri, D. J. Fleet, R. Kiros, and S. Fidler. VSE++: improved visual-semantic embeddings. *CoRR*, abs/1707.05612, 2017. [1](#)
- [4] A. Frome, G. Corrado, J. Shlens, S. Bengio, J. Dean, M. Ranzato, and T. Mikolov. Devise: A deep visual-semantic embedding model. In *Neural Information Processing Systems (NIPS)*, 2013. [1](#)
- [5] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. [2](#)
- [6] D. Kiela and L. Bottou. Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 36–45, Doha, Qatar, October 2014. Association for Computational Linguistics. [1](#)
- [7] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *CoRR*, abs/1411.2539, 2014. [1](#)
- [8] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *CoRR*, abs/1411.2539, 2014. [1](#)
- [9] K. Kurach, S. Gelly, M. Jastrzebski, P. Häusser, O. Teytaud, D. Vincent, and O. Bousquet. Better text understanding through image-to-text transfer. *CoRR*, abs/1705.08386, 2017. [1](#)
- [10] A. Lazaridou, N. T. Pham, and M. Baroni. Combining language and vision with a multimodal skip-gram model. In R. Mihalcea, J. Y. Chai, and A. Sarkar, editors, *HLT-NAACL*, pages 153–163. The Association for Computational Linguistics, 2015. [1](#)
- [11] H. Ling and S. Fidler. Teaching machines to describe images via natural language feedback. *CoRR*, abs/1706.00130, 2017. [1](#)
- [12] J. Mao, J. Xu, K. Jing, and A. L. Yuille. Training and evaluating multimodal word embeddings with large-scale web annotated images. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Advances in Neural Information Processing Systems 29*, pages 442–450. Curran Associates, Inc., 2016. [1](#)
- [13] M. Pagliardini, P. Gupta, and M. Jaggi. Unsupervised learning of sentence embeddings using compositional n-gram features. *CoRR*, abs/1703.02507, 2017. [1](#)
- [14] A. Parikh, O. Täckström, D. Das, and J. Uszkoreit. A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255, Austin, Texas, November 2016. Association for Computational Linguistics. [2](#)
- [15] Y. Peng, J. Qi, and Y. Yuan. Modality-specific cross-modal similarity measurement with recurrent attention network. *CoRR*, abs/1708.04776, 2017. [1](#)
- [16] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, 2014. [2](#)
- [17] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier. Collecting image annotations using amazon’s mechanical turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 139–147, Los Angeles, June 2010. Association for Computational Linguistics. [2](#)
- [18] C. Silberer and M. Lapata. Learning grounded meaning representations with autoencoders. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 721–732, Baltimore, Maryland, June 2014. Association for Computational Linguistics. [1](#)
- [19] R. Socher, A. Karpathy, Q. Le, C. Manning, and A. Ng. Grounded compositional semantics for finding and describing images with sentences. *Transactions of the Association for Computational Linguistics*, 2:207–218, 2014. [1](#)
- [20] Y. H. Tsai, L. Huang, and R. Salakhutdinov. Learning robust visual-semantic embeddings. *CoRR*, abs/1703.05908, 2017. [1](#)