

Semantic Relatedness Based Re-ranker for Text Spotting

Ahmed Sabir¹ Francesc Moreno-Noguer² Lluís Padró¹

¹ Universitat Politècnica de Catalunya, TALP Research Center, Barcelona, Spain

² Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain

Abstract

Applications such as textual entailment, plagiarism detection or document clustering rely on the notion of *semantic similarity*, and are usually approached with dimension reduction techniques like LDA or with embedding-based neural approaches. We present a scenario where semantic similarity is not enough, and we devise a neural approach to learn *semantic relatedness*. The scenario is *text spotting in the wild*, where a text in an image (*e.g.* street sign, advertisement or bus destination) must be identified and recognized. Our goal is to improve the performance of vision systems by leveraging semantic information. Our rationale is that the text to be spotted is often related to the image context in which it appears (word pairs such as *Delta-airplane*, or *quarters-parking* are not *similar*, but are clearly *related*). We show how learning a word-to-word or word-to-sentence relatedness score can improve the performance of text spotting systems up to 2.9 points, outperforming other measures in a benchmark dataset.

1 Introduction

Deep learning has been successful in tasks related to deciding whether two short pieces of text refer to the same topic, *e.g.* semantic textual similarity (Cer et al., 2018), textual entailment (Parikh et al., 2016) or answer ranking for Q&A (Severyn and Moschitti, 2015).

However, other tasks require a broader perspective to decide whether two text fragments are *related* more than whether they are *similar*. In this work, we describe one of such tasks, and we re-train some of the existing sentence similarity approaches to learn this semantic relatedness. We also present a new Deep Neural Network (DNN) that outperforms existing approaches when applied to this particular scenario.

The task we tackle is *Text Spotting*, which is the problem of recognizing text that appears in unrestricted images (a.k.a. *text in the wild*) such as traffic signs, commercial ads, or shop names. Current state-of-the-art results on this task are far from those of OCR systems with simple backgrounds.

Existing approaches to Text Spotting usually divide the problem in two fundamental tasks: 1) *text detection*, consisting of selecting the image regions likely to contain texts, and 2) *text recognition*, that converts the images within these bounding boxes into a readable string. In this work, we focus on the recognition stage, aiming to prove that semantic relatedness between the image context and the recognized text can be useful to boost the system performance. We use existing pre-trained architectures for Text Recognition, and add a shallow deep-network that performs a post-processing operation to re-rank the proposed candidate texts. In particular, we re-rank the candidates using their semantic relatedness score with other visual information extracted from the image (*e.g.* objects, scenario, image caption). Extensive evaluation shows that our approach consistently improves other semantic similarity methods.

2 Text Hypothesis Extraction

We use two pre-trained Text Spotting baselines to extract k text hypotheses. The first baseline is a CNN (Jaderberg et al., 2016) with fixed lexicon based recognition, able to recognize words in a predefined 90K-word dictionary. Second, we use an LSTM architecture with visual attention model (Ghosh et al., 2017) that generates the final output words as probable character sequences, without relying on any lexicon. Both models are trained on a synthetic dataset (Jaderberg et al., 2014). The output of both models is a Softmax score for each of the k candidate words.

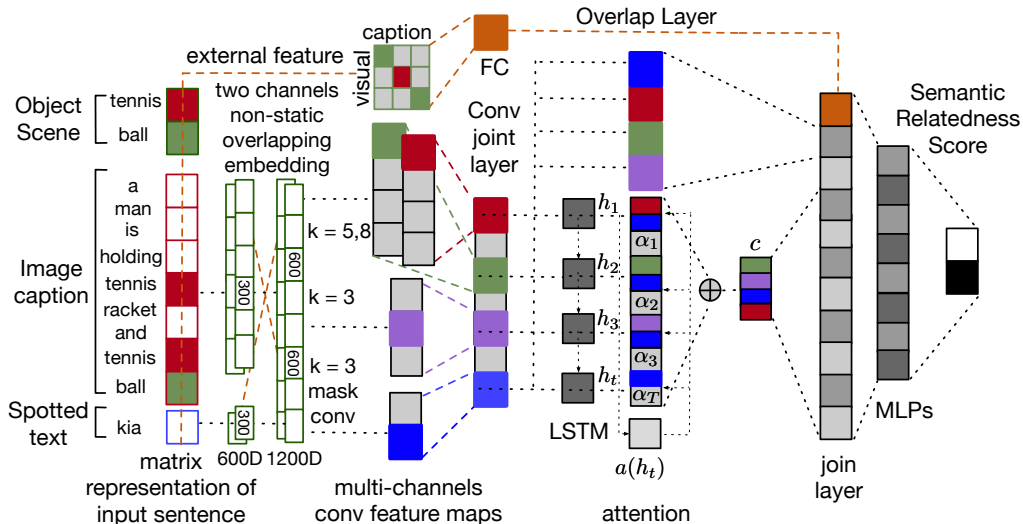


Figure 1: Overview of the system pipeline, an end-to-end post-process scores the semantic relatedness between a candidate word and the context in the image (objects, scenarios, natural language descriptions, ...)

3 Learning Semantic Relatedness for Text Spotting

To learn the semantic relatedness between the visual context information and the candidate word we introduce a multi-channel convolutional LSTM with an attention mechanism. The network is fed with the candidate word plus several words describing the image visual context (object and places labels, and descriptive captions)¹, and is trained to produce a relatedness score between the candidate word and the context.

Our architecture is inspired by (Severyn and Moschitti, 2015), that proposed CNN-based re-rankers for Q&A. Our network consists of two subnetworks, each with 4-channels with kernel sizes $k = (3, 3, 5, 8)$, and overlap layer, as shown in Figure 1. We next describe the main components:

Multi-Channel Convolution: The first subnetwork consists of only convolution kernels, and aims to extract n-gram or keyword features from the caption sequence.

The convolution is applied over a sequence to extract n-gram features from different positions. Let $x \in \mathbb{R}^{s \times d}$ be the sentence matrix, where s is the sentence length, and d the dimension of the i -th word in the sentence. Let also denote by $c \in \mathbb{R}^{k \times d}$ the kernel for the convolution operation. For each i -th position in the sentence, w_i is the concatenation of k consecutive words, i.e.,

$w_i = [x_i \oplus x_{i+1} \oplus \dots \oplus x_{i+k-1}]$. Our architecture uses multiple such kernels to generate feature maps m . The feature map for each window vector w_i can be written as:

$$m_i = f(w_i \circ c + b) \quad (1)$$

where \circ is element-wise multiplication, f is non-linear function, in our case we apply Relu function (Nair and Hinton, 2010), and b is a bias. For j kernels, the generated j feature maps can be arranged as feature representation for each window W_i as: $W = [m_1 \oplus m_2 \oplus \dots \oplus m_j]$. Each row W_i of $W \in \mathbb{R}^{(s-k+1) \times j}$ is the new generated feature from the j -th kernel for the window vector at position i . The new generated feature (window representations) are then fed into the joint layer and LSTM as shown in Figure 1.

Multi-Channel Convolution-LSTM: Following C-LSTM (Zhou et al., 2015) we forward the output of the CNN layers into an LSTM, which captures the long term dependencies over the features. We further introduce an attention mechanism to capture the most important features from that sequence. The advantage of this attention is that the model learns the sequence without relying on the temporal order. We describe in more detail the attention mechanism below.

Also, following (Zhou et al., 2015), we do not use a pooling operation after the convolution feature map. Pooling layer is usually applied after the convolution layer to extract the most important features in the sequence. However, the output of our Convolutional-LSTM model is fed into

¹All this visual context information is automatically generated using off-the-shelf existing modules (see section 4).

an LSTM (Hochreiter and Schmidhuber, 1997) to learn the extracted sequence, and pooling layer would break that sequence via downsampling to a selected feature. In short, LSTM is specialized in learning sequence data, and pooling operation would break such a sequence order. On the other hand, for the Multi-Channel Convolution model we also learn the extracted word sequence n -gram directly and without feature selection, pooling operation.

Attention Mechanism: Attention-based models have shown promising results on various NLP tasks (Bahdanau et al., 2014). Such mechanism learns to focus on a specific part of the input (e.g. a relevant word in a sentence). We apply an attention mechanism (Raffel and Ellis, 2015) via an LSTM that captures the temporal dependencies in the sequence. This attention uses a Feed Forward Neural Network (FFN) attention function:

$$e_t = \tanh(h_t W_a) v_a^T \quad (2)$$

where W_a is the attention of the hidden weight matrix and v_a is the output vector. As shown in Fig. 1 the vector c is computed as a weighted average of h_t , given by α (defined below). The attention mechanism is used to produce a single vector c for the complete sequence as follows:

$$e_t = a(h_t), \alpha_t = \frac{\exp(e_t)}{\sum_{k=1}^T \exp(e_k)}, c = \sum_{t=1}^T \alpha_t h_t \quad (3)$$

where T is the total number of steps and α_t is the computed weight of each time step t for each state h_t , a is a learnable function that depends only on h_t . Since this attention computes the average over time, it discards the temporal order, which is ideal for learning semantic relations between words. By doing this, the attention gives higher weights to more important words in the sentence without relying on sequence order.

Overlap Layer: The overlap layer is just a frequency count dictionary to compute overlap information of the inputs. The idea is to give more weight to the most frequent visual element, specially when it is observed by more than one visual classifier. The dictionary output is a fully connected layer.

Finally, we merge all subnetworks into a joint layer that is fed to a loss function which calculates the semantic relatedness between both inputs. We call the combined model Fusion Dual Convolution-LSTM-Attention (FDCLSTM_{AT}).

Since we have only one candidate word at a time, we apply a convolution with *masking* in the candidate word side (first channel). In this case, simply zero-padding the sequence has a negative impact on the learning stability of the network. We concatenate the CNN outputs with the additional feature into MLP layers, and finally a sigmoid layer performing binary classification. We trained the model with a binary cross-entropy loss (l) where the target value (in $[0, 1]$) is the semantic relatedness between the word and the visual. Instead of restricting ourselves to a simple similarity function, we let the network learn the margin between the two classes –i.e. the degree of similarity. For this, we increase the depth of network after the MLPs merge layer with more fully connected layers. The network is trained using Nesterov-accelerated Adam (Nadam) (Dozat, 2016) as it yields better results (specially in cases such as word vectors/neural language modelling) than other optimizers using only classical momentum (ADAM). We apply batch normalization (BN) (Ioffe and Szegedy, 2015) after each convolution, and between each MLPs layer. We omitted the BN after the convolution for the model without attention (FDCLSTM), as BN deteriorated the performance. Additionally, we consider 70% dropout (Srivastava et al., 2014) between each MLPs for regularization purposes.

4 Dataset and Visual Context Extraction

We evaluate the performance of the proposed approach on the COCO-text (Veit et al., 2016). This dataset is based on Microsoft COCO (Lin et al., 2014) (Common Objects in Context), which consists of 63,686 images, and 173,589 text instances (annotations of the images). This dataset does not include any visual context information thus we used out-of-the-box object (He et al., 2016) and place (Zhou et al., 2014) classifiers and tuned a caption generator (Vinyals et al., 2015) on the same dataset to extract contextual information from each image, as seen in Figure 2.

5 Related Work and Contribution

Understanding the visual environment around the text is very important for scene understanding. This has been recently explored by a relatively reduced number of works. Zhu et al. (2016) shows that the visual context could be beneficial for text detection. This work uses a 14 classes pixel clas-

Table 1: Best results after re-ranking using different re-ranker, and different values for k -best hypotheses extracted from the baseline output (%). In addition, To test our re-ranker with MRR we fixed k CNN $_{k=8}$ LSTM $_{k=4}$

Model	CNN					LSTM			
	full	dict	list	k	MRR	full	list	k	MRR
Baseline (BL)	full: 19.7 dict: 56.0					full: 17.9			
BL+ Glove (Pennington et al., 2014)	22.0	62.5	75.8	7	44.5	19.1	75.3	4	78.8
BL+C-LSTM (Zhou et al., 2015)	21.4	61.0	71.3	8	45.6	18.9	74.7	4	80.7
BL+CNN-RNN (Wang et al., 2016)	21.7	61.8	73.3	8	44.5	19.5	77.1	4	80.9
BL+MVCNN (Yin and Schütze, 2016)	21.3	60.6	71.9	8	44.2	19.2	75.8	4	78.8
BL+Attentive LSTM (Tan et al., 2016)	21.9	62.4	74.0	8	45.7	19.1	71.4	5	80.2
BL+fasttext (Joulin et al., 2017)	21.9	62.2	75.4	7	44.6	19.4	76.1	4	80.3
BL+InferSent (Conneau et al., 2017)	22.0	62.5	75.8	7	44.5	19.4	76.7	4	79.7
BL+USE-T (Cer et al., 2018)	22.0	62.5	78.3	6	44.7	19.2	75.8	4	79.5
BL+TWE (Sabir et al., 2018)	22.2	63.0	76.3	7	44.7	19.5	76.7	4	80.2
BL+FDCLSTM (ours)	22.3	63.3	75.1	8	45.0	20.2	67.9	9	79.8
BL+FDCLSTM _{AT} (ours)	22.4	63.7	75.5	8	45.9	20.1	67.6	9	81.8
BL+FDCLSTM _{lexicon} (ours)	22.6	64.3	76.3	8	45.1	19.4	76.4	4	78.8
BL+FDCLSTM _{AT+lexicon} (ours)	22.6	64.3	76.3	8	45.1	19.7	77.8	4	80.4

sifier to extract *context features* from the image, such as *tree, river, wall*, to then assist scene text detection. Kang et al. (2017) employs topic modeling to learn the correlation between visual context and the spotted text in social media. The meta-data associated with each image (e.g tags, comments and titles) is then used as context to enhance recognition accuracy. Karaoglu et al. (2017) takes advantage of text and visual context for logo retrieval problem. Most recently, Prasad and Wai Kin Kong (2018) use object information (limited to 42 predefined object classes) surrounding the spotted text to guide text detection. They propose two sub-networks to learn the relation between text and object class (e.g. relations such as *car-plate* or *sign board-digit*).

Unlike these methods, our approach uses direct visual context from the image where the text appears, and does not rely on any extra resource such as human labeled meta-data (Kang et al., 2017) nor limits the context object classes (Prasad and Wai Kin Kong, 2018). In addition, our approach is easy to train and can be used as a drop-in complement for any text-spotting algorithm that outputs a ranking of word hypotheses.

6 Experiments and Results

In the following we use different *similarity* or *relatedness* scorers to reorder the k -best hypothesis produced by an off-the-shelf state-of-the-art text spotting system. We experimented extracting k -best hypotheses for $k = 1 \dots 10$.

We use two pre-trained deep models: a CNN (Jaderberg et al., 2016) and an LSTM (Ghosh et al., 2017) as baselines (BL) to extract the ini-

tial list of word hypotheses.

The CNN baseline uses a closed lexicon; therefore, it cannot recognize any word outside its 90K-word dictionary. Table 1 presents four different accuracy metrics for this case: 1) **full** columns correspond to the accuracy on the whole dataset. 2) **dict** columns correspond to the accuracy over the cases where the target word is among the 90K-words of the CNN dictionary (which correspond to 43.3% of the whole dataset. 3) **list** columns report the accuracy over the cases where the right word was among the k -best produced by the baseline. 4) **MRR** Mean Reciprocal Rank (MRR), which is computed as follows: $MRR = \frac{1}{|Q|} \sum_{k=1}^{|Q|} \frac{1}{rank_k}$, where rank k is the position of the correct answer in the hypotheses list proposed by the baseline.

Comparing with sentence level model: We compare the results of our encoder with several state-of-the-art sentence encoders, tuned or trained on the same dataset. We use cosine to compute the similarity between the caption and the candidate word. Word-to-sentence representations are computed with: Universal Sentence Encoder with the Transformer USE-T (Cer et al., 2018), and In-fersent (Conneau et al., 2017) with glove (Pennington et al., 2014). The rest of the systems in Table 1 are trained in the same conditions that our model with glove initialization with dual-channel overlapping² non-static pre-trained embedding on the same dataset. Our model FDCLSTM without attention achieves a better result in the case of the second baseline LSTM that full of false-positives and short words. The advantage of the attention mechanism is the ability to integrate information

²To increase the overlap embedding dimensions



Figure 2: Examples of candidate re-ranking using object (c1), place (c2), and caption (c3) information. The *delta-airliner* relation which frequently co-occurs in training data is captured by the *overlap* layers. The *12-football* pair shows a relatedness between sports and numbers.

over time, and it allows the model to refer to specific points in the sequence when computing its output. However, in this case, the attention attends the wrong context, as there are many words have no correlation or do not correspond to actual words. On the other hand, USE-T seems to require a shorter hypothesis list to get top performance when the right word is in the hypothesis list.

Comparing with word level model: We also compare our result with current state-of-the-art word embeddings trained on a large general text using *glove* and *fasttext*. The word model used only object and place information, and ignored the caption. Our proposed models achieve better performance than our TWE previous model (Sabir et al., 2018), that trained a word embedding (Mikolov et al., 2013) from scratch on the same task.

Similarity to probabilities: After computing the cosine similarity we need to convert that score to probabilities. As we proposed in previous work (Sabir et al., 2018) we obtain the final probability combining (Blok et al., 2003) the similarity score, the probability of the detected context (provided by the object/place classifier), and the probability of the candidate word (estimated from a 5M token corpus) (Lison and Tiedemann, 2016).

Effect of Unigram Probabilities: Ghosh et al. (2017) showed the utility of a language model (LM) when the data is too small for a DNN, obtaining significant improvements. Thus, we introduce a basic model of unigram probabilities with Out-of-vocabulary (OOV) words smoothing. The model is applied at the end, to re-rank out false positive short words, and has the main goal of re-ranking out less probable word overranked by the deep model. As seen in Table 1, the introduction

of this unigram lexicon produces the best results.

Human performance: To estimate an upper bound for the results, we picked 33 random pictures from the test dataset and had 16 human subjects try to select the right word among the top $k = 5$ candidates produced by the baseline text spotting system. Our proposed model performance on the same images was 57%. Average human performance was 63% (highest 87%, lowest 39%).

7 Conclusion

In this work, we propose a simple deep learning architecture to learn semantic relatedness between word-to-word and word-to-sentence pairs, and show how it outperforms other semantic similarity scorers when used to re-rank candidate answers in the Text Spotting problem.

In the future, we plan using the same approach to tackle similar problems, including lexical selection in Machine Translation, or word sense disambiguation (Lala and Specia, 2018). We believe our approach could also be useful in multimodal machine translation, where an image caption must be translated using not only the text but also the image content (Barrault et al., 2018). Tasks that lie at the intersection of computer vision and NLP, such as the challenges posed in the new BreakingNews dataset (popularity prediction, automatic text illustration) could also benefit from our results (Ramisa et al., 2018).

Acknowledgments

We would like to thank José Fonollosa, Marta Ruiz Costa-Jussà and the anonymous reviewers for discussion and feedback. This work was supported by the KASP Scholarship Program and by the MINECO project HuMoUR TIN2017-90086-R.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Loïc Barrault, Fethi Bougares, Lucia Specia, Chirraag Lala, Desmond Elliott, and Stella Frank. 2018. Findings of the third shared task on multimodal machine translation. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 304–323.
- Sergey Blok, Douglas Medin, and Daniel Osherson. 2003. Probability from similarity. In *AAAI Spring Symposium on Logical Formalization of Commonsense Reasoning*.
- Daniel Cer, Yinfei Yang, Sheng-yi Kong, Nan Hua, Nicole Limtiaco, Rhomni St John, Noah Constant, Mario Guajardo-Cespedes, Steve Yuan, Chris Tar, et al. 2018. Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Timothy Dozat. 2016. Incorporating nesterov momentum into adam.
- Suman K Ghosh, Ernest Valveny, and Andrew D Bagdanov. 2017. Visual attention models for scene text recognition. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, volume 1, pages 943–948. IEEE.
- 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*.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.
- Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2014. Synthetic data and artificial neural networks for natural scene text recognition. *arXiv preprint arXiv:1406.2227*.
- Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2016. Reading text in the wild with convolutional neural networks. *International Journal of Computer Vision*, 116(1):1–20.
- Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Association for Computational Linguistics.
- Chulmoo Kang, Gunhee Kim, and Suk I Yoo. 2017. Detection and recognition of text embedded in on-line images via neural context models. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Sezer Karaoglu, Ran Tao, Jan C van Gemert, and Theo Gevers. 2017. Con-text: Text detection for fine-grained object classification. *IEEE Transactions on Image Processing*.
- Chirraag Lala and Lucia Specia. 2018. Multimodal lexical translation. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*.
- 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*. Springer.
- Pierre Lison and Jörg Tiedemann. 2016. Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*.
- Vinod Nair and Geoffrey E Hinton. 2010. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th international conference on machine learning (ICML-10)*.
- Ankur Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2249–2255.
- 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)*.
- Shitala Prasad and Adams Wai Kin Kong. 2018. Using object information for spotting text. In *The European Conference on Computer Vision (ECCV)*.
- Colin Raffel and Daniel PW Ellis. 2015. Feed-forward networks with attention can solve some long-term memory problems. *arXiv preprint arXiv:1512.08756*.

- Arnau Ramisa, Fei Yan, Francesc Moreno-Noguer, and Krystian Mikolajczyk. 2018. Breakingnews: Article annotation by image and text processing. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(5):1072–1085.
- Ahmed Sabir, Francesc Moreno-Noguer, and Lluís Padró. 2018. Visual re-ranking with natural language understanding for text spotting. In *Asian Conference on Computer Vision*, pages 68–82. Springer.
- Aliaksei Severyn and Alessandro Moschitti. 2015. Learning to rank short text pairs with convolutional deep neural networks. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*. ACM.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Ming Tan, Cicero Dos Santos, Bing Xiang, and Bowen Zhou. 2016. Improved representation learning for question answer matching. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 464–473.
- Andreas Veit, Tomas Matera, Lukas Neumann, Jiri Matas, and Serge Belongie. 2016. Coco-text: Dataset and benchmark for text detection and recognition in natural images. *arXiv preprint arXiv:1601.07140*.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164.
- Xingyou Wang, Weijie Jiang, and Zhiyong Luo. 2016. Combination of convolutional and recurrent neural network for sentiment analysis of short texts. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*.
- Wenpeng Yin and Hinrich Schütze. 2016. Multichannel variable-size convolution for sentence classification. *arXiv preprint arXiv:1603.04513*.
- Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. 2014. Learning deep features for scene recognition using places database. In *Advances in neural information processing systems*.
- Chunting Zhou, Chonglin Sun, Zhiyuan Liu, and Francis Lau. 2015. A c-lstm neural network for text classification. *arXiv preprint arXiv:1511.08630*.
- Anna Zhu, Renwu Gao, and Seiichi Uchida. 2016. Could scene context be beneficial for scene text detection? *Pattern Recognition*.