
Recurrent Instance Segmentation using Sequences of Referring Expressions

Alba Maria Herrera-Palacio
Universitat Politècnica de Catalunya

Carles Ventura
Universitat Oberta de Catalunya
cventuraroy@uoc.edu

Carina Silberer
Universitat Pompeu Fabra
carina.silberer@upf.edu

Ionut-Teodor Sorodoc
Universitat Pompeu Fabra
ionutteodor.sorodoc@upf.edu

Gemma Boleda
Universitat Pompeu Fabra
gemma.boleda@upf.edu

Xavier Giro-i-Nieto
Universitat Politècnica de Catalunya
xavier.giro@upc.edu

Abstract

The goal of this work is to segment the objects in an image that are referred to by a sequence of linguistic descriptions (referring expressions). We propose a deep neural network with recurrent layers that output a sequence of binary masks, one for each referring expression provided by the user. The recurrent layers in the architecture allow the model to condition each predicted mask on the previous ones, from a spatial perspective within the same image. Our multimodal approach uses off-the-shelf architectures to encode both the image and the referring expressions. The visual branch provides a tensor of pixel embeddings that are concatenated with the phrase embeddings produced by a language encoder. Our experiments on the RefCOCO dataset for still images indicate how the proposed architecture successfully exploits the sequences of referring expressions to solve a pixel-wise task of instance segmentation.

1 Introduction

In this work, we tackle object instance segmentation with natural language expressions, a challenging problem with implications in the fields of computer vision and natural language processing. The goal is to segment the referent, i.e., the target object referred to by a referring expression, in an image. For instance, given the image in Figure 1(a) and the referring expression “left woman in blue”, the model needs to output the mask for the relevant person (Figure 1(d)). Instance segmentation with referring expressions can be understood as an extension of semantic instance segmentation, where a binary mask and a categorical label are assigned to each object in an image (see comparison in Figure 1). Humans use referring expressions to talk about objects in the world; therefore, the ability to ground referring expressions in images can be very useful in human-computer interaction scenarios, too.

Work in this area [5, 7, 11, 13] separately represents the linguistic expression and the input image, typically using recurrent neural networks (RNN) and convolutional neural networks (CNN), respectively. Afterwards, in order to obtain a pixel-wise segmentation mask, both representations are combined and further processed. In the case of multiple referring expressions over the same image, each of them is processed separately. More details about related work are contained in the supplementary material. We focus on the novel scenario in which a user does not provide a single referring expression, but a



Figure 1: Comparison between different segmentation tasks: (b) object segmentation, (c) object instance segmentation and (d) segmentation from natural language expressions.

sequence of them, one for each referent. For each expression in the sequence, our model predicts a visual grounding conditioned by not only the current reference, but also the previous ones. In addition, our model is end-to-end trainable and does not require any visual post-processing as in MAttNet [14], which was based on the Mask R-CNN computer vision model for instance segmentation [3]. Mask R-CNN, and other similar solutions, predicts a large amount of instances which are later filtered.

The proposed architecture consists of: (i) a vision encoder, which extracts visual features of a frame, (ii) a language encoder, which adds linguistic information to the model by using a pre-trained natural language processing model to extract language features for the referring expressions (phrases), and (iii) a recurrent segment decoder, which uses the image and phrase embeddings from the vision and language encoders, respectively, to generate the pixel-level masks of the target objects.

2 Method

We propose an end-to-end trainable deep neural network to recurrently segment target objects indicated by linguistic referring expressions (RE). The proposed architecture is depicted in Figure 2. The visual encoder and decoder are inspired by RSIS [10], a deep neural network for object instance segmentation. The language embeddings for the referring expressions are obtained from the BERT encoder [2]. The pixel and phrase embeddings are concatenated and fed to a binary mask visual decoder. Given a sequence of linguistic referring expressions, the recurrent nature of the mask decoder allows to condition the current prediction on the previous ones.

The BERT [2] encoding, represented at the top branch of Figure 2, is used without any fine-tuning. We use the base model of 12 encoder layers (transformer blocks), 768 hidden units and 12 attention heads.¹ Given a referring expression, BERT outputs a set of contextualized embeddings which comprises the hidden states of each encoder layer of each word (token). We average the last hidden layer of each token producing a single 768 length vector for each referring expression. To avoid memory problems while training the model and to balance the dimensions of the language and visual embeddings, we reduce the dimensionality of the textual embeddings to 64 with principal component analysis (PCA) [9].

The visual encoding and decoding schemes were adopted from the RSIS [10] model for semantic instance segmentation. The input image is encoded with a ResNet-101 [4] model pre-trained on ImageNet [1]. The ResNet architecture is truncated at the last convolutional layer, thus removing the last two layers (pooling layer and classification layer). In contrast to the language branch, the image encoder was finetuned for the task. The output of each convolutional block is used as an image feature, which provides a set of visual features at different resolutions, as shown in dark blue at the left of Figure 2.

The input to the mask decoder for a given referent consists of a set of multi-resolution pixel embeddings obtained by the visual encoder, and the phrase embedding provided by the language encoder. Consequently, visual features are shared among all the referents for the same image, and the output of the decoder is a sequence of object segmentation predictions, one for each referent.

¹Publicly available as bert-base-cased model at <https://github.com/huggingface/pytorch-transformers>.

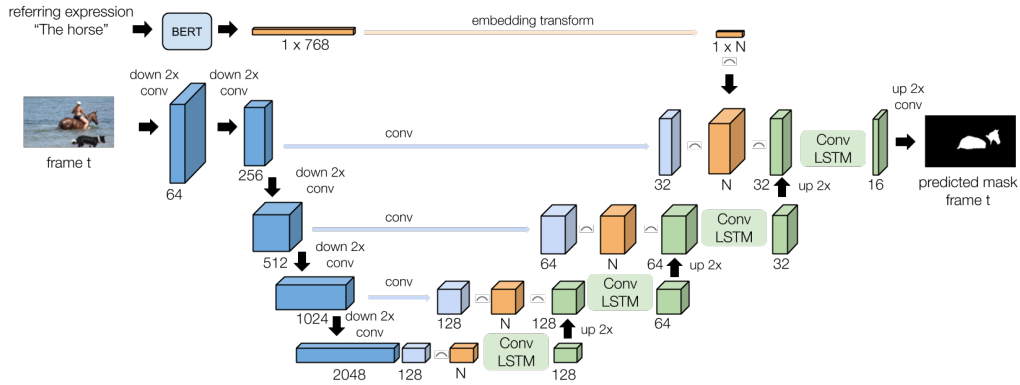


Figure 2: Our proposed recurrent architecture for recurrent instance segmentation with linguistic referring expressions. The figure illustrates a single forward pass, predicting only the mask of one instance for an image.

The mask decoder is an extension of the multi-resolution one proposed in RSIS [10]. In order to keep the inherent spatial information in the visual features when segmenting an instance, for each resolution, we concatenate the corresponding language embedding to each feature map along the channels’ dimensions (depth) of the visual tensors. This allows every pixel embedding to receive the whole representation of the language information. The ConvLSTM [12] layers used in the decoder allows to condition the predicted masks with those masks predicted for previously presented referring expressions over the same image.

Similarly to [10], the cost function is defined as the soft Intersection over Union score between the predicted mask and the ground truth mask for a given referent. Since we do sequential processing, during training we use as ground truth the mask corresponding to the referring expression being processed at each timestep.

3 Experiments

The experiments show how the introduction of the referring expression encoder successfully conditions the mask to predict, and that the order of the phrases within the input sequences affects the performance of the model. The experiments have been performed on the RefCOCO dataset [15], a dataset with 142K referring expressions for 50K objects in 20K images from MSCOCO [8], that is, where the target objects are of 80 common categories. More details on the dataset and model training are contained in the supplementary material.

We validate the performance of the referring expression branch by comparing the results with the baseline case of not using referring expressions. In this case, RSIS is used to generate a fixed-length sequence of instance masks. The length of the sequence is always larger than the amount of reference phrases associated with the image, avoiding to penalize the segmentation of objects for which no referring expression is presented to the model. Instead of forcing a specific order when matching the predicted masks and ground truth masks, the Hungarian algorithm [6] carries out an optimal assignment between them using the soft Intersection over Union score as cost function. The results presented in Table 1 show four different configurations in terms of referent order and batch size, with (our multimodal model) or without (RSIS, i.e. the visual model) referring expressions. Our solution consistently outperforms RSIS, even when RSIS is completely free to generate its masks in any order. These results show that the linguistic phrases are successfully used to identify the right target instance, and that, in addition, the quality of the masks actually improves over the language-free task addressed by RSIS.




Figure 3: Qualitative results: (a) original image, (b) ground truth, (c) segmentation result, (d) and (e) pixel-wise predictions (red pixels are false negatives, blue true positives and green false positives).

We also investigated the effect of the order within the sequence of referring expressions used to train the model. We considered two options: (i) by area, (ii) randomly. The results in Table 1 indicate that the best strategy is to randomly feed the referents and use small batch sizes. The fact that our best result is obtained for the smallest batch size (16) and a random ordering may indicate that our model overfits and that further reducing the amount of parameters to learn may even increase the performance. If we focus on the batch size 32 with the referring expression, we can also observe that the random configuration almost doubles the performance with respect to training with objects sorted by area. These results highlight the importance to randomize the training expressions to avoid learning undesirable data biases.

Figure 3 shows some qualitative results generated by our network. The depicted results are among the good predictions of the algorithm and show how our model can distinguish between different instances of the same class.

Finally, Figure 4 depicts how the order of the segmented objects is consistent with the order of the referring expressions. By reversing the order of the phrases, the order of the generated masks is also reversed, which shows that the model has learnt how to associate REs with visual objects. Note that the generated masks are not exactly the same. This indicates that the model indeed conditions its segmentation decisions on its predictions for previous REs. Note how the generated masks are not exactly the same, another evidence that suggests that the order affects the segmentation results. Further qualitative results are contained in the supplementary material.

4 Conclusions

This work has proposed a solution for visual object segmentation by adding a new linguistic branch to the RSIS deep neural architecture. The concatenation of the phrase embeddings of the referring expression to each pixel embedding of the RSIS decoder has the potential to successfully condition the predicted mask to the desired object. The recurrent nature of the decoder allows to process sequences

Table 1: Results on RefCOCO with and without referring expressions.

Referent order	Batch size	Referring expression	Instance IoU \uparrow			Overall IoU \uparrow		
			val	testA	testB	val	testA	testB
By area	128		21.82	25.56	18.86	18.48	21.27	16.48
	128	✓	26.08	29.63	22.81	23.67	26.47	21.13
	32		21.68	23.50	19.67	19.42	21.02	17.94
	32	✓	26.12	28.66	23.82	23.88	25.81	22.23
Random	128		20.36	22.70	15.78	17.65	19.32	15.22
	128	✓	27.54	31.45	24.39	24.75	27.76	22.26
	32		20.13	23.13	19.04	17.77	19.83	17.24
	32	✓	39.79	45.31	34.04	35.70	40.28	31.28
	16	✓	42.66	47.48	37.51	36.95	41.42	32.72




Figure 4: Changing referring expressions order: (a) original image, (b) ground truth, (c) segmentation result (original order), and (d) segmentation results (inverse order).

of referring phrases and condition the output based on the previous predictions. The proposed architecture is trained end-to-end and avoids the additional computation required by post-processing steps such as non-maximum suppression or ranking of object proposals.

Acknowledgements

UPC: This work has been developed in the framework of project TEC2016-75976-R, funded by the Spanish Ministerio de Economía y Competitividad and the European Regional Development Fund (ERDF), and the Industrial Doctorate 2017-DI-011 funded by the Government of Catalonia. We gratefully acknowledge the support of NVIDIA Corporation with the donation of some of the GPUs used for this work.

UOC: This work has been partially supported by the Ministerio de Economía, Industria y Competitividad (Spain), under the Grant Ref. RTI2018-095232-B-C22.

UPF: This project has also received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 715154), and from the Ramón y Cajal programme (grant RYC-2015-18907). We gratefully acknowledge the support of NVIDIA Corporation with the donation of GPUs used for this research, and the computer resources at CTE-POWER and the technical support provided by Barcelona Supercomputing Center (RES-FI-2018-3-0034). This paper reflects the authors’ view only, and the EU is not responsible for any use that may be made of the information it contains.

References

- [1] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*, 2009.
- [2] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [3] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017.
- [4] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [5] Ronghang Hu, Marcus Rohrbach, and Trevor Darrell. Segmentation from natural language expressions. In *European Conference on Computer Vision*, pages 108–124. Springer, 2016.
- [6] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955.

- [7] Ruiyu Li, Kaican Li, Yi-Chun Kuo, Michelle Shu, Xiaojuan Qi, Xiaoyong Shen, and Jiaya Jia. Referring image segmentation via recurrent refinement networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2018.
- [8] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [9] Karl Pearson. Liii. on lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 2(11):559–572, 1901.
- [10] Amaia Salvador, Miriam Bellver, Victor Campos, Manel Baradad, Ferran Marques, Jordi Torres, and Xavier Giro-i Nieto. Recurrent neural networks for semantic instance segmentation. *arXiv preprint arXiv:1712.00617*, 2017.
- [11] Hengcan Shi, Hongliang Li, Fanman Meng, and Qingbo Wu. Key-word-aware network for referring expression image segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 38–54, 2018.
- [12] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. Convolutional lstm network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*, pages 802–810, 2015.
- [13] Linwei Ye, Mrigank Rochan, Zhi Liu, and Yang Wang. Cross-modal self-attention network for referring image segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10502–10511, 2019.
- [14] Licheng Yu, Zhe Lin, Xiaohui Shen, Jimei Yang, Xin Lu, Mohit Bansal, and Tamara L Berg. Mattnet: Modular attention network for referring expression comprehension. In *CVPR*, 2018.
- [15] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *European Conference on Computer Vision*, pages 69–85. Springer, 2016.