

Detecting and Grouping Identical Objects for Region Proposal and Classification

Abbeloos, W.; Caccamo, S.; Ataer-Cansizoglu, E.; Taguchi, Y.; Feng, C.; Lee, T.-Y.

TR2017-099 July 2017

Abstract

Often multiple instances of an object occur in the same scene, for example in a warehouse. Unsupervised multiinstance object discovery algorithms are able to detect and identify such objects. We use such an algorithm to provide object proposals to a CNN-based classifier. This results in fewer regions to evaluate, compared to traditional region proposal algorithms. Additionally, it enables using the joint probability of multiple instances of an object, resulting in improved classification accuracy. The proposed technique can also split a single class into multiple sub-classes corresponding to the different object types, enabling hierarchical classification.

CVPR Workshop on Deep Learning for Robotic Vision

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

Detecting and Grouping Identical Objects for Region Proposal and Classification

Wim Abbeloos *
KU Leuven, Belgium

Sergio Caccamo *
KTH Royal Institute of Technology, Sweden

Esra Ataer-Cansizoglu, Yuichi Taguchi, Chen Feng and Teng-Yok Lee
Mitsubishi Electric Research Labs (MERL), USA

Abstract

Often multiple instances of an object occur in the same scene, for example in a warehouse. Unsupervised multi-instance object discovery algorithms are able to detect and identify such objects. We use such an algorithm to provide object proposals to a CNN-based classifier. This results in fewer regions to evaluate, compared to traditional region proposal algorithms. Additionally, it enables using the joint probability of multiple instances of an object, resulting in improved classification accuracy. The proposed technique can also split a single class into multiple sub-classes corresponding to the different object types, enabling hierarchical classification.

1. Introduction

Recent years have seen tremendous progress in object detection and classification based on convolutional neural networks (CNNs). A standard pipeline is using an object proposal method and classifying each region using a CNN [2] as in Figure 1 (top). While general object proposal methods work well for general scenes, some scenes such as industrial and warehouse environments have more structures and the objects show less variations in appearance.

In this work, we exploit such structures by using a multi-instance object discovery algorithm [1] that is able to discover, localize, and identify object instances that occur in a scene multiple times. The algorithm uses RGB-D images as the input and searches for patterns of local features that occur in multiple objects. The output of the algorithm is clusters of the features, each corresponding to one instance of an object. The clusters corresponding to multiple instances of the same object are associated with each other. As shown in Figure 1 (bottom), bounding boxes of the discovered ob-

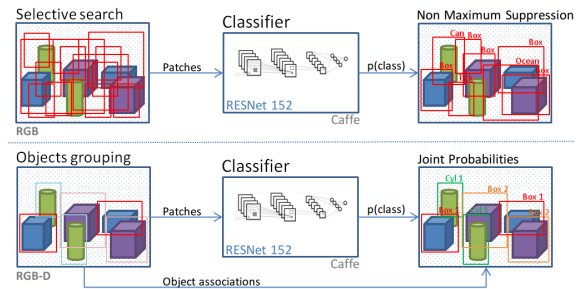


Figure 1. (Top) A typical object detection pipeline utilizing a CNN. (Bottom) The proposed pipeline using an alternative object discovery method for region proposal.

jects are used as the input to a classification network. Our object proposal method also provides the identity of the object, allowing to accumulate the probabilities and further improve the classification accuracy.

2. Experiments and Results

A dataset consisting of ten scenes with six cereal boxes on a table was captured using an ASUS XTion RGB-D camera. There were three different object instances, two instances of each type (Figure 2). We made sure the same faces of the objects of the same type were visible so that they could be discovered, identified, and localized properly.

For our method the bounding boxes were obtained by taking each feature cluster expanded by 60 percent to ensure they encapsulate the entire object. The resulting regions were cropped and resized to have 640 pixels on the longest side. These images were then classified using ResNet-152 [3], pre-trained on ImageNet (1000 classes). This resulted in a vector of 1000 class probabilities.

For the baseline method the bounding boxes returned by selective search [4] were resized and classified using the same ResNet-152. An extra non-maximum suppression step was used with a maximum overlap threshold of 0.6. A detection was considered correct if the intersection over union was at least 0.25. The ground truth poses of the boxes

*Denotes equal authorship. WA and SC performed the work during their internship at MERL. Corresponding author email: cansizoglu@merl.com

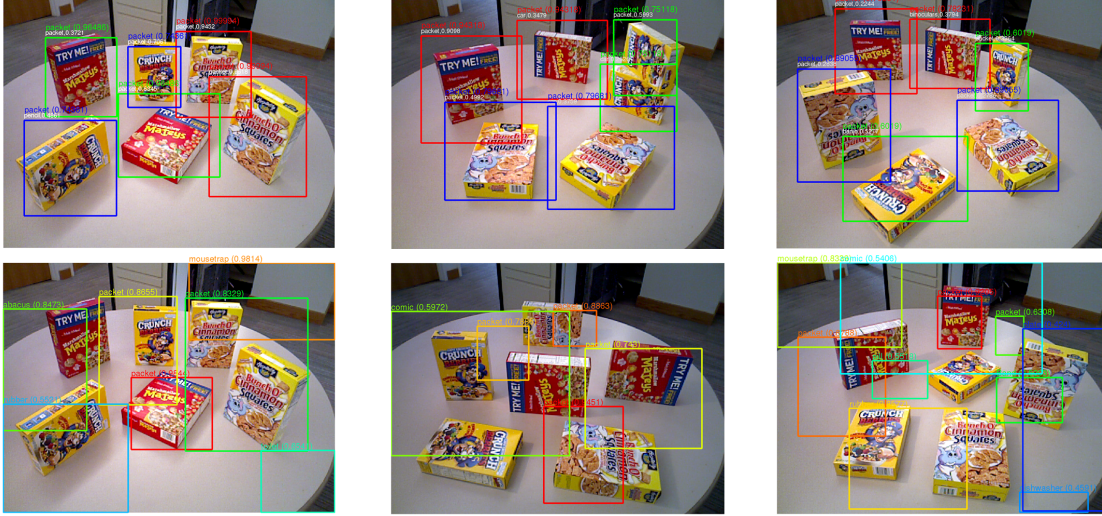


Figure 2. Detection and classification results using our method (top) and the baseline method (bottom). For our method the bounding box colors indicate object association. We show both the individual (white) and joint (in color) probabilities. For the baseline method we only show the regions with probability larger then 0.4 for better visualization.

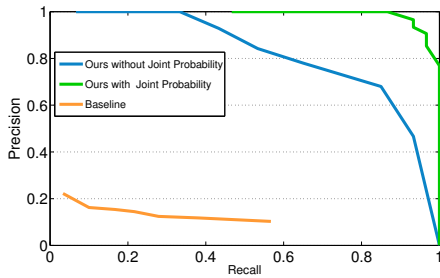


Figure 3. The precision-recall curves for our method without (blue) and with (green) use of the joint probability. The red curve is the baseline method.

were manually annotated.

The sum of the probabilities of the categories “packet” and “comic” was used as the probability of the object being a cereal box. To evaluate the classification results, the precision and recall were calculated as

$$P = \frac{\sum_{i=1}^C tp_i}{\sum_{i=1}^C (tp_i + fp_i)}, \quad R = \frac{\sum_{i=1}^C tp_i}{\sum_{i=1}^C (tp_i + fn_i)}, \quad (1)$$

with tp_i , fp_i , and fn_i the true positives, false positives, and false negatives, respectively.

Figure 3 shows the precision-recall curves. For our method we compare the result where we consider the objects as independent (without computing the joint probability), and that where we compute the joint probability of all instances of the same object. The joint probability was computed by multiplying the class probability vectors of these instances and re-normalizing them, which increased the probabilities of correct classifications and improved the precision-recall performance.

Figure 2 visually shows some of the scenes with the detected and classified objects. As can be seen we have a much smaller number of detected regions, while the baseline method tends to generate many false positives. The average number of proposed regions by our method was six, while that of the baseline method was 400. After non-maximum suppression the average number of remaining regions was 20.

3. Discussion

Our object detection framework proposes object regions based on discovering multiple instances of identical, feature rich objects. While these are strong assumptions, several real-world robotic scenarios satisfy these conditions, which would benefit from our method. Currently our evaluation was performed using a limited number of simplified scenes. We plan to extend the evaluation using a larger number of more complex scenes.

References

- [1] W. Abbeloos, S. Caccamo, E. Ataer-Cansizoglu, Y. Taguchi, and Y. Domaec. 3D object discovery and modeling using single RGB-D images containing multiple object instances. Submitted to IROS, 2017. 1
- [2] R. Girshick. Fast R-CNN. In *Proc. IEEE Int’l Conf. Computer Vision (ICCV)*, pages 1440–1448, 2015. 1
- [3] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 1
- [4] J. R. Uijlings, K. E. Van De Sande, T. Gevers, and A. W. Smeulders. Selective search for object recognition. *Int’l J. Computer Vision*, 104(2):154–171, 2013. 1