Content-Aware Image Resizing for Faster Object Detection on Aerial Imagery

Anouk Visser, Arnoud Visser, Jan van Gemert

University of Amsterdam Delft University of Technology

Abstract. To combat poaching or perform game counts nature conservationists need to inspect areas that are very large and hard to reach by car or foot. Recently, nature conservationists have been able to inspect these areas more easily by using UAVs equipped with cameras. Despite of the ease with which these systems can be deployed, the recorded imagery still needs to be analyzed manually. Automatic object detection algorithms could greatly reduce the time spent looking for the object of interest, benefiting the conservation work.

State-of-the-art object detection algorithms such as R-CNN [1] rely heavily on object proposals that are used to provide a speedup in the object detection pipeline. Object proposal methods such as Selective Search [2] or Edge Boxes [3] have proven to work well on popular datasets such as PASCAL VOC [4] and ImageNet [5]. However, the resolution of the images in these datasets are much lower than the image resolution needed for nature conservation tasks. Object proposal methods are significantly slower when applied to high resolution images, which affects the detection rate. To maintain a good object detection rate, we apply a content-aware image resizing method that resizes the image without compromising on the content.

Figure 1 shows the original image and a version that was reduced to 25% of its original size. Even though the majority of the pixels have been removed, the objects (cows) are still clearly visible. Figure 2 compares our method to regular resizing by comparing the recall at different image sizes. The figure shows that as we reduce the image more, the recall drops much faster when using regular resizing than when using our content-aware image resizing method. We show that we can use our content-aware image resizing to speed up object detection on aerial imagery without significantly impacting detection performance.

References

1. Girshick, R.: Fast r-cnn. arXiv preprint arXiv:1504.08083 (2015)
2. Uijlings, J.R., van de Sande, K.E., Gevers, T., Smeulders, A.W.: Selective search for object recognition. International journal of computer vision **104**(2) (2013) 154–171
3. Zitnick, C.L., Dollár, P.: Edge boxes: Locating object proposals from edges. In: Computer Vision–ECCV 2014. Springer (2014) 391–405
4. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2007 (VOC2007) Results. http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.htm

(a)

(b)

Fig. 1: Image from the Verschoor Aerial Cow Dataset 1a Original image. 1b Image reduced to 25% of its original size. Even though the majority of the pixels have been removed, the cows are still clearly visible.

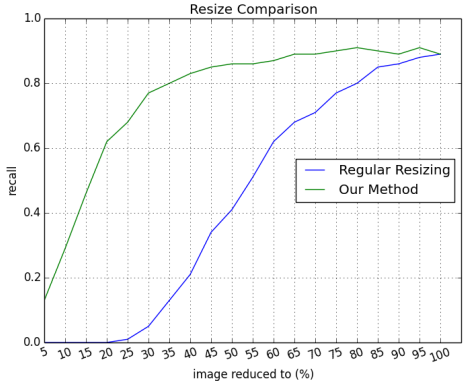

Fig. 2: Recall of Selective Search on the Verschoor Aerial Cow Dataset after reducing the image size. *Regular Resizing* shows the recall over different scale factors after simply scaling down the images, *Our Method* shows the recall when the image was resized by our content-aware image resizing method. We see that as we reduce the image more, the recall drops much faster when using regular resizing than when using our content-aware image resizing method.

5. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. International Journal of Computer Vision (2014) 1–42