Boundary-aware Instance Segmentation: Supplementary Material

Zeeshan Hayder[1,2], Xuming He[2,1]
[1]Australian National University & [2]Data61/CSIRO

Mathieu Salzmann[3]
[3]CVLab, EPFL, Switzerland

1. Boundary-aware Instance Segmentation: Detailed Comparisons with MNC on Pascal VOC 2012

We now focus on the detailed comparison with MNC [1], which is the method most closely related to ours and achieves the best performance after us. In Table 1, we provide the detailed evaluation over all the classes of the Pascal VOC 2012 dataset [2] using IoU thresholds of 0.5 and 0.7, respectively. Note that our method outperforms this baseline for most classes, with a particularly large margin for an IoU of 0.7.

2. Object Mask Network: Comparison with MNC and DeepMask for Different Object Sizes

We further analyze the quality of the segmentation proposals for different sizes of objects. Following the criteria of [3], all object instances are categorized into small, medium, or large. The average recall is computed for the top 100 proposals using the mask level intersection over union. In Table 2, we provide the results of this evaluation on the Pascal VOC 2012 validation set and a comparison with the MNC [1], DeepMask [4] and SharpMask [5] proposals. While SharpMask achieves the best performance on small objects, our approach yields competitive results on medium objects and clearly outperforms all these baselines on large ones.

3. Qualitative Comparison with MNC

In Fig. 1, we compare our boundary-aware instance segmentation results with those obtained by MNC. These results again evidence the better quality of our instance-level semantic segmentations.

VOC 2012	IoU	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	mAP
MNC-new	0.5	80.1	67.6	71.7	49.6	47.4	80.1	70.2	86.6	37.1	62.9	37.0	83.1	75.2	75.6	76.1	40.3	70.1	47.8	76.4	65.6	65.0
BAIS (ours)	0.5	80.1	68.3	72.2	47.0	49.2	79.8	70.8	86.0	39.0	65.8	39.6	84.8	73.7	75.9	76.9	41.1	70.9	47.7	78.6	66.2	**65.7**
MNC-new	0.7	58.1	33.9	49.5	32.9	34.6	71.6	53.7	75.4	17.9	48.2	18.5	66.9	43.4	48.9	48.1	19.7	49.6	34.1	64.6	55.5	46.2
BAIS (ours)	0.7	60.2	39.7	52.1	31.4	36.0	71.8	56.8	75.7	20.4	51.5	21.8	69.3	47.7	52.8	51.0	21.3	51.6	35.3	63.9	55.7	**48.3**

Table 1. **Detailed Evaluation on the Pascal VOC 2012 validation set.** We report results with mask level IoU thresholds of 0.5 and 0.7. Note that our method outperforms MNC [1] for most classes and achieves an improvement of 2.1% over MNC at a 0.7% IoU threshold.

PASCAL VOC 2012	AR@Small	AR@Medium	AR@Large
DeepMask [4]	24.6	43.9	50.5
SharpMask [5]	**25.8**	**47.6**	55.3
MNC+score	19.4	44.0	60.7
OMN (ours)	20.3	45.9	**63.4**

Table 2. **Object size analysis on the Pascal VOC 2012 validation set.** We compare our method with state-of-the-art segmentation proposal baselines according to the criteria of [3]. The AR for small, medium and large objects were computed for 100 proposals. Note that our object mask network outperforms the state-of-the-art baselines for large objects by a significant margin.

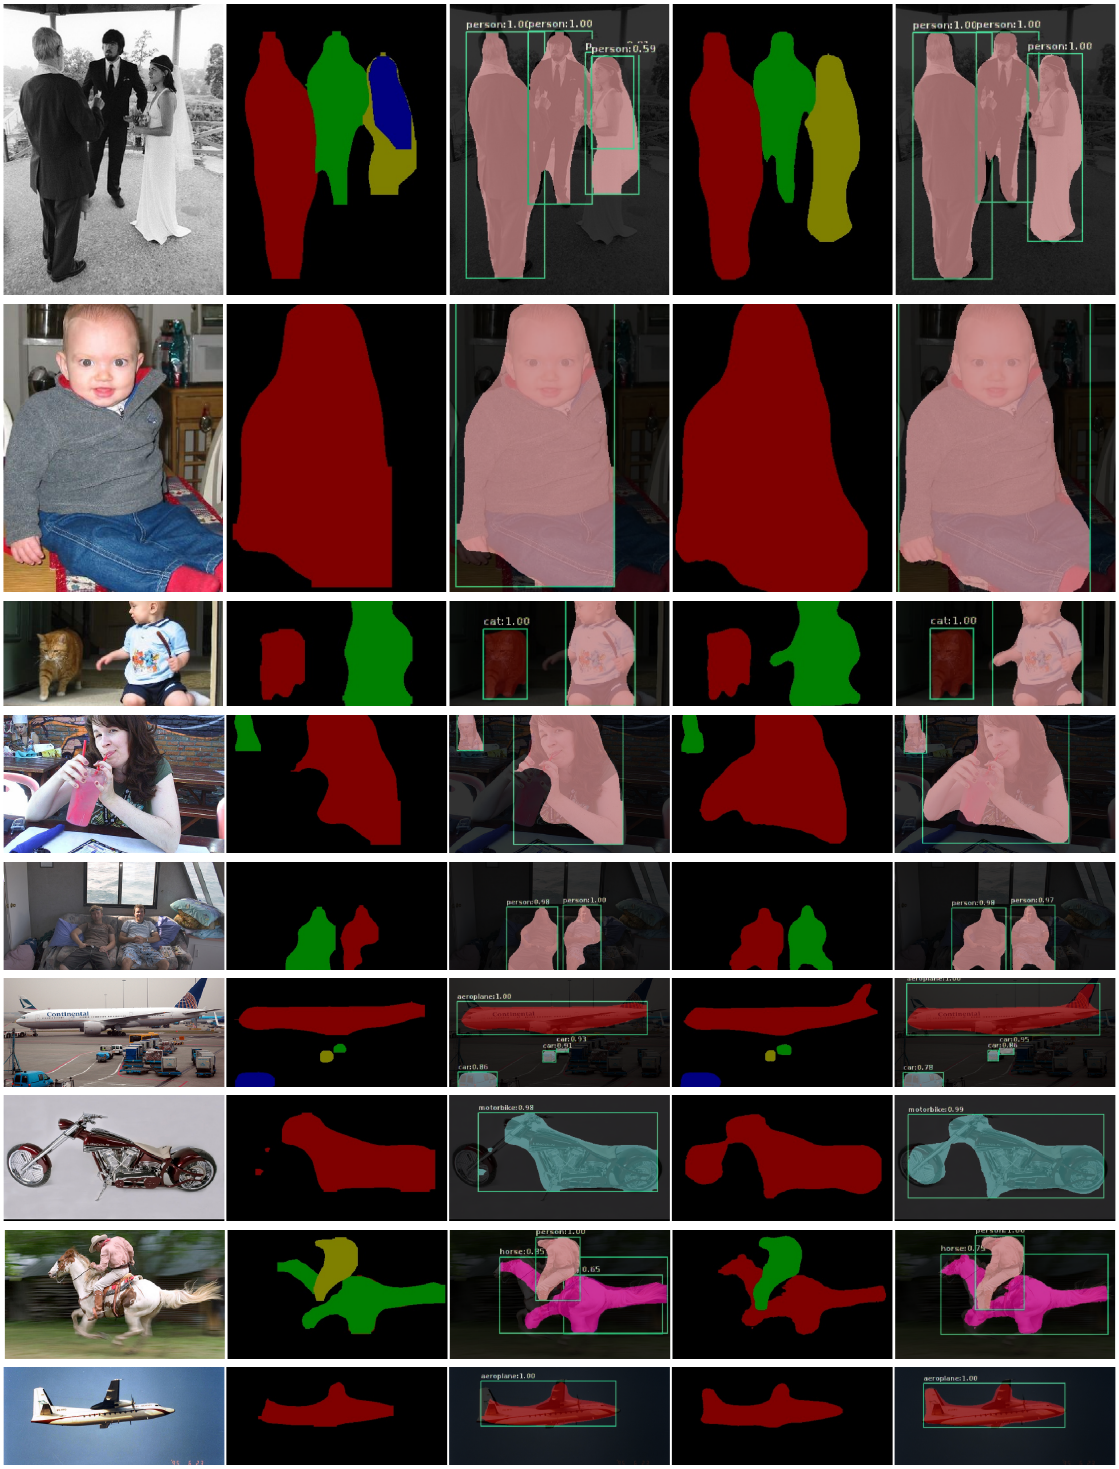

Figure 1. **Qualitative results on Pascal VOC 2012.** From left to right, we show: the original image, the segmentations of MNC, these results projected on the image, our segmentations, our results projected on the image. Note the better quality of our results.

References

[1] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *CVPR*, 2016. 1

[2] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html. 1

[3] T. Lin, M. Maire, S. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. 1

[4] P. O. Pinheiro, R. Collobert, and P. Dollr. Learning to segment object candidates. In *NIPS*, 2015. 1

[5] P. O. Pinheiro, T.-Y. Lin, R. Collobert, and P. Dollr. Learning to refine object segments. In *ECCV*, 2016. 1