Detection and Segmentation of 2D Curved Reflection Symmetric Structures: Supplementary Material

Ching L. Teo cteo@cs.umd.edu
Cornelia Fermüller fer@umiacs.umd.edu
Yiannis Aloimonos yiannis@cs.umd.edu
Computer Vision Lab, University of Maryland, College Park, MD 20742, USA

Abstract

We provide full parameters and additional qualitative results of our curved symmetry detector (SRFSym) and symmetry-constrained segmentation (SymSegGC) over all four datasets: 1) SYMMAX-300 (SYMSEG-300), 2) BSD-Parts, 3) WHD and 4) NY-roads, described in the main paper.

1. Parameters

The same set of parameters for training and evaluating our approach are used over all four datasets. Tables 1 and 2 summarize the parameters used for training the SRF-based curved symmetry detector (SRFSym) and the symmetry-constrained segmentation (SymSegGC) respectively. We also used the best parameters for other approaches reported in the paper. This is obtained by training over different training subsets from the four datasets considered. For RTree [2], we used the published results and dataset available online[1].

The key evaluation parameters that generate the Precision-Recall (P-R) curves per dataset are summarized in Table 3. The same evaluation parameters were used for all approaches for fairness. Notations and descriptions follow the main paper.

Notation	Description	Value
-	Number of feature orientations	8 [0 to π]
-	Number of feature scales	4 ([0.1, 0.3, 0.5, 0.75] of image diagonal)
-	Number of feature channels	6 [L^*, a^*, b^*, textons, spectral, Gabor]
N	Patch size	16
-	Number of [positive/negative] training samples per dataset	[10^5/10^5]
$\|\mathcal{B}\|$	Size of structured labels	150
K	Number of trees	16
h_d	Maximum tree depth	64
-	Minimum value of M_i	10^{-10}
-	Minimum length of A (pixels)	5
-	Minimum symmetry response, $A(r)$ per pixel $r \in A$	0.01

Table 1. Parameters used in SRFSym.

2. Additional Results

We first show additional results from SYMMAX-300 that compare segmentation with and without the symmetry prior in Fig. 1.

Figs. 2, 3, 4 and 5 show additional curved symmetry detections and segmentation of curved symmetric structures for example test images from SYMMAX-300 (SYMSEG-300), BSD-Parts, WHD and NY-roads respectively and compare them with DefDiscs-SP [1] for the first three datasets and RTree-ES [2] for NY-roads.

[1] https://documents.epfl.ch/users/a/as/asironi/www/results_Aerial_precomputed.zip

Notation	Description	Value
-	Edge response threshold for I_e	0.03
σ	Smoothing factor in boundary pairwise term, V_{pq}	0.32
β	Penalty for symmetrical neighbors in $S_{pp'}$	0.006
ρ_b	Ballooning cost in B_{pq}	-0.03

Table 2. Parameters used in `SymSegGC`.

Experiment type	Dataset	# of recall values	Matching distance threshold	Min segment size (pixels)	Min Intersect-over-Union (IoU) score
Curved Symmetry Detection	SYMMAX-300 [3]	50	0.01% of of image diagonal	-	-
	NY-roads [2]	50	2 pixels	-	-
Symmetry-constrained Segmentation	SYMSEG-300	-	-	200	0.4
	BSD-Parts [1]	-	-	200	0.4
	WHD [1]	-	-	50	0.4
	NY-roads	20	-	0	0 (after application of exclusion zone)†

Table 3. Evaluation parameters. †For NY-roads, precision is computed as the intersection of the thresholded segmentation with the groundtruth divided by the thresholded segmentation, after applying an *exclusion* zone of 0.4% of the groundtruth radius. See [2] for details.

References

[1] T. S. H. Lee, S. Fidler, and S. Dickinson. Detecting curved symmetric parts using a deformable disc model. In *ICCV*, pages 1753–1760, 2013.

[2] A. Sironi, V. Lepetit, and P. Fua. Multiscale centerline detection by learning a scale-space distance transform. In *CVPR*, pages 2697–2704, 2014.

[3] S. Tsogkas and I. Kokkinos. Learning-based symmetry detection in natural images. In *ECCV*, pages 41–54. 2012.

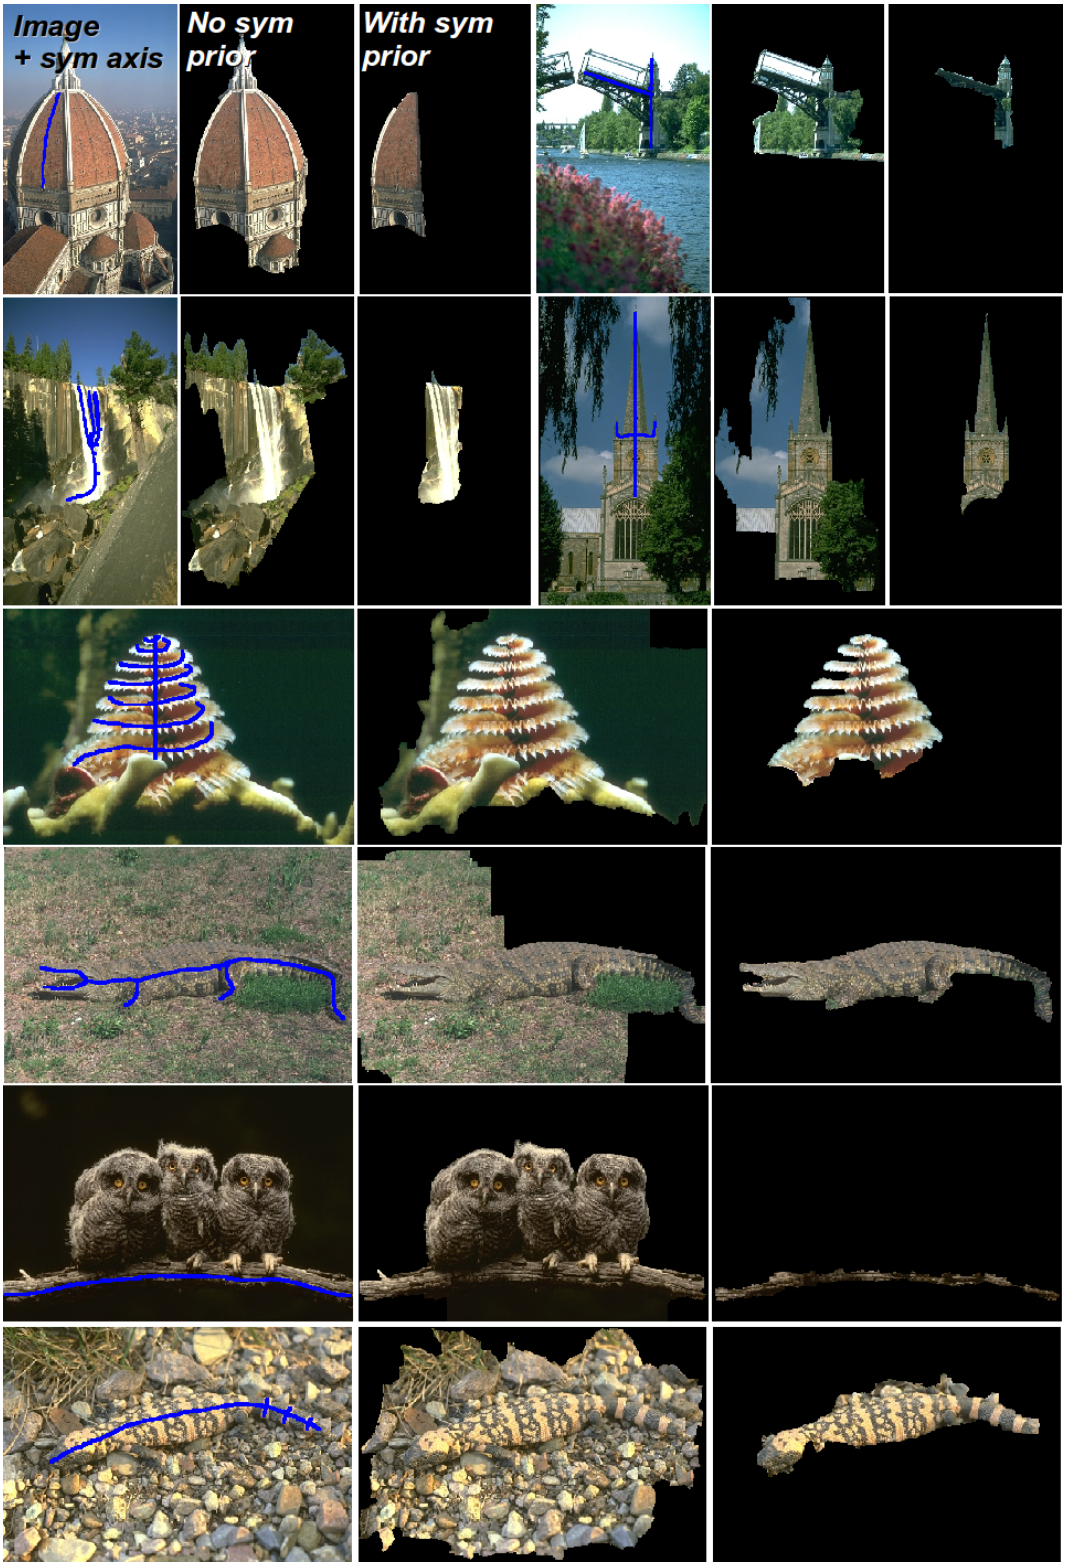

Figure 1. Additional results illustrating the contribution of the symmetry prior. 8 results (L-R): Input image with example curved symmetry axis (blue), segmentation without symmetry prior, segmentation with symmetry prior.

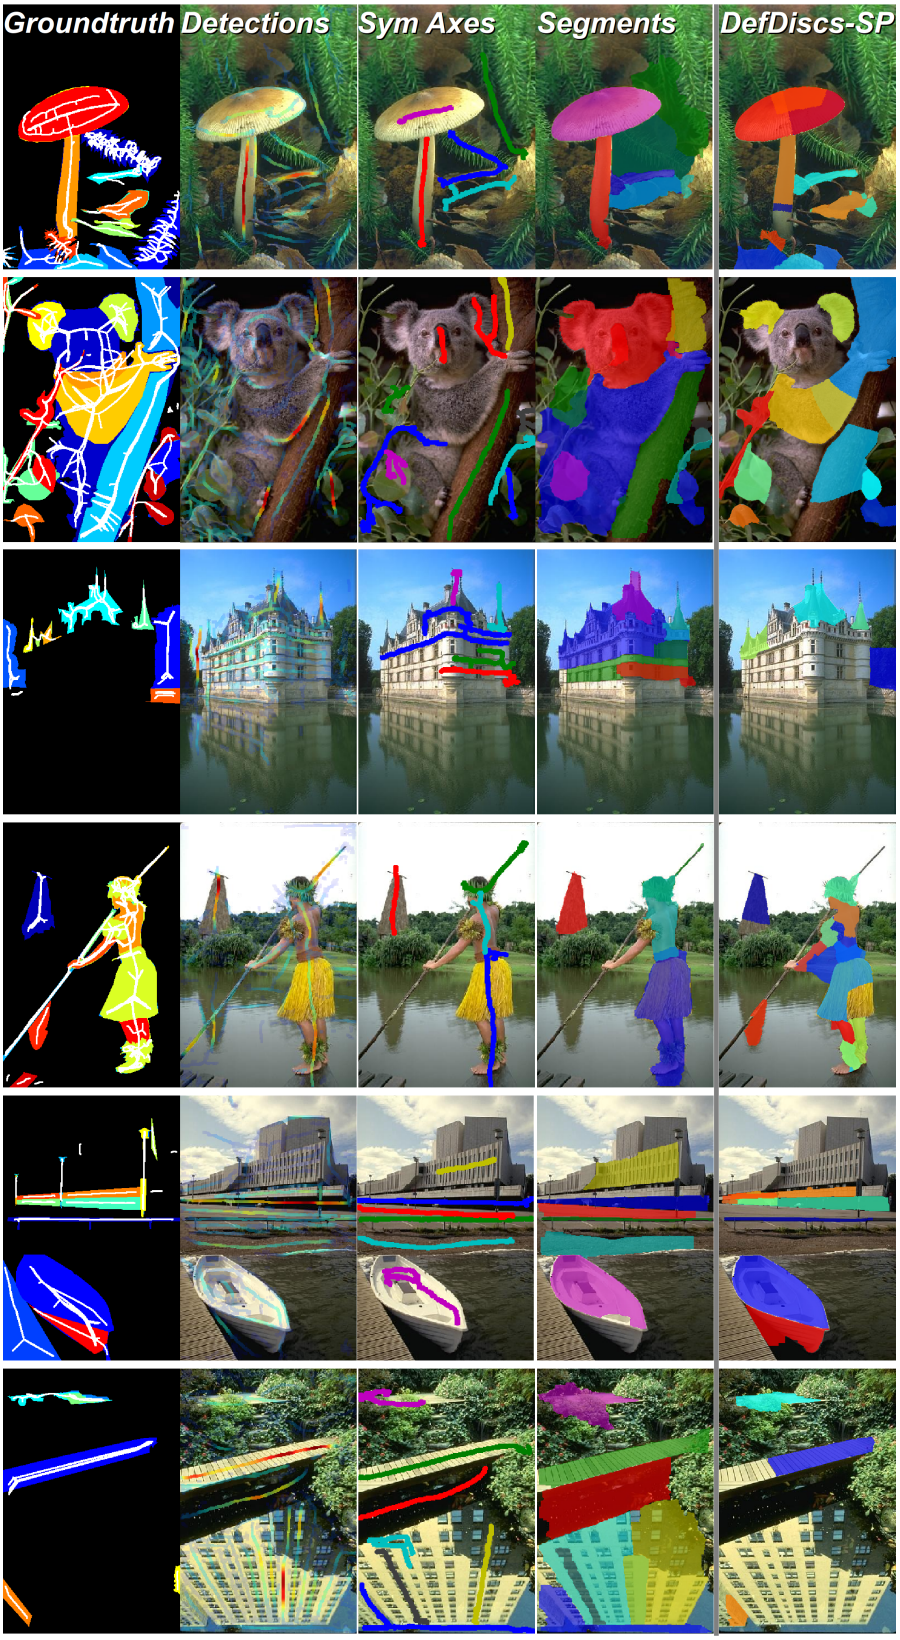

Figure 2. Additional results for SYMMAX-300 (SYMSEG-300). 6 results (L-R): Groundtruth symmetry axes (white) and regions, `SRFSym` detections, Selected symmetry axes, `SymSegGC` segmentations color coded with symmetry axes, `DefDiscs-SP` segmentations.

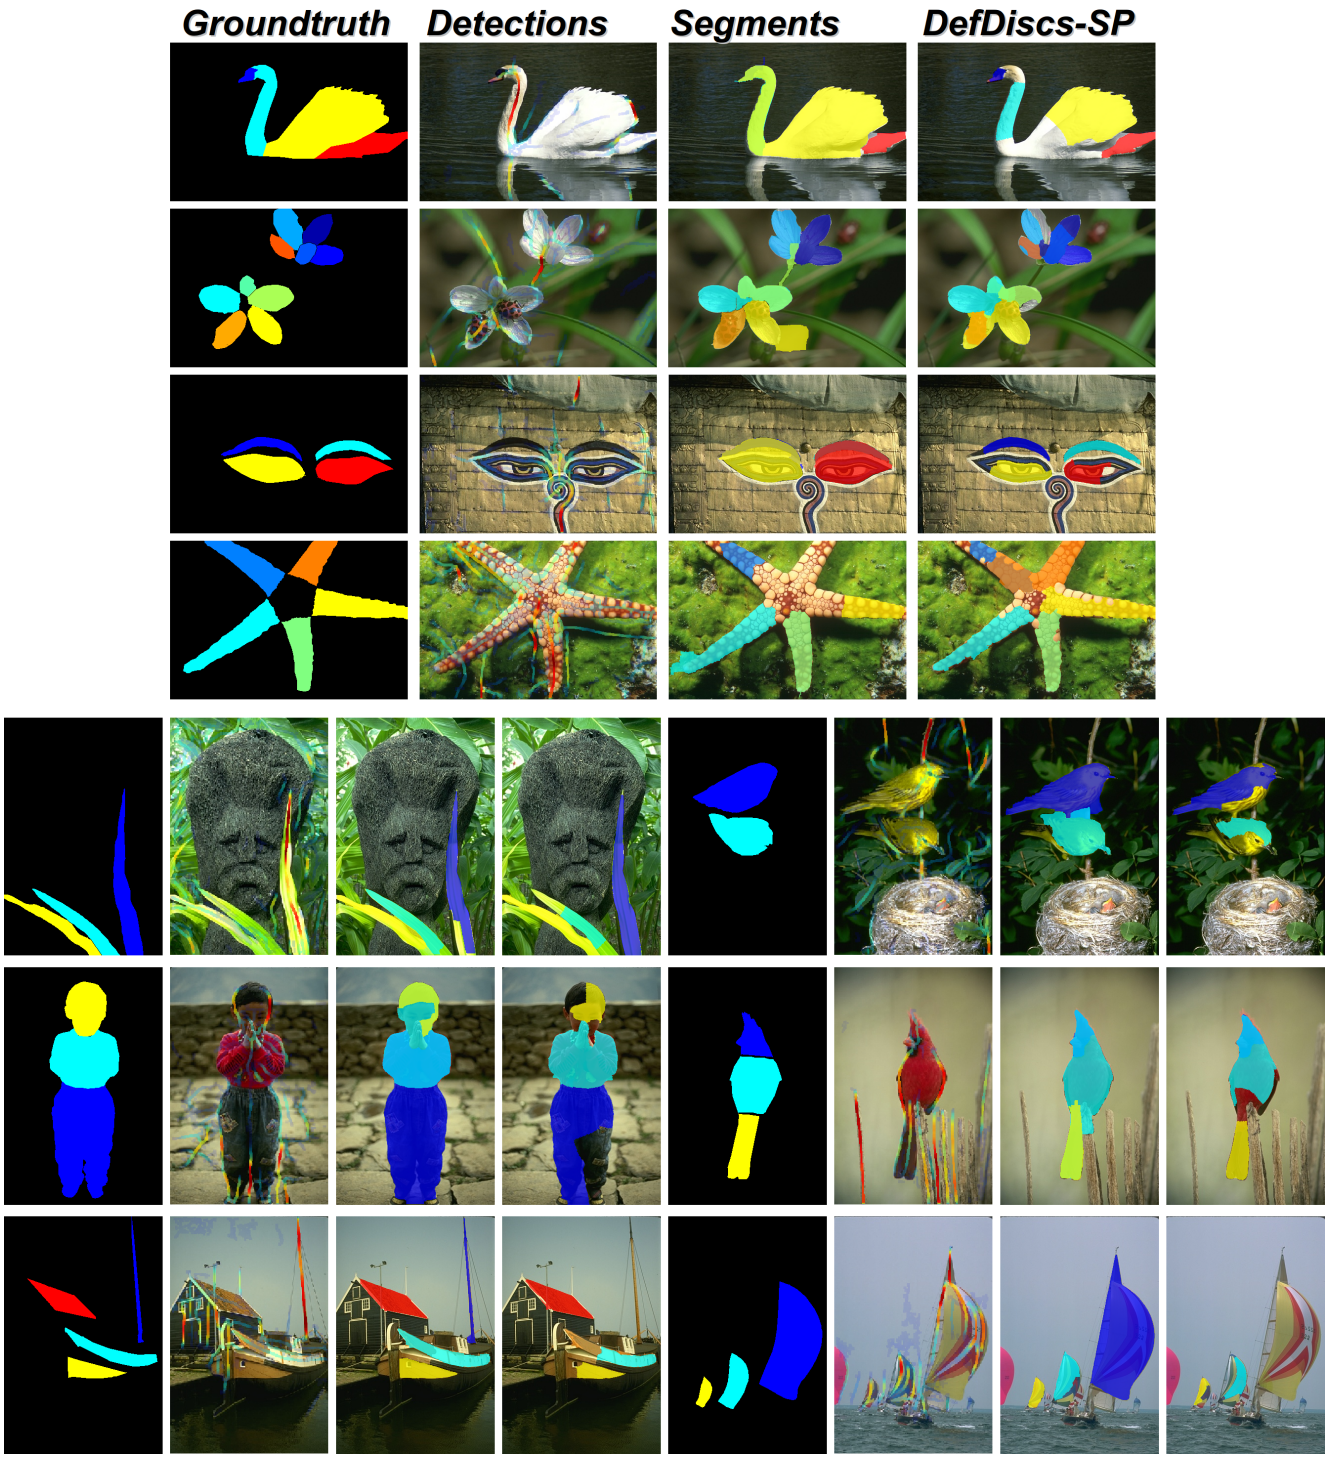

Figure 3. Additional results for BSD-Parts. 10 results (L-R): Groundtruth symmetrical regions, `SRFSym` detections, `SymSegGC` segmentations, `DefDiscs-SP` segmentations. We show here segmentations that have significant overlap with the groundtruth, and are color coded accordingly.

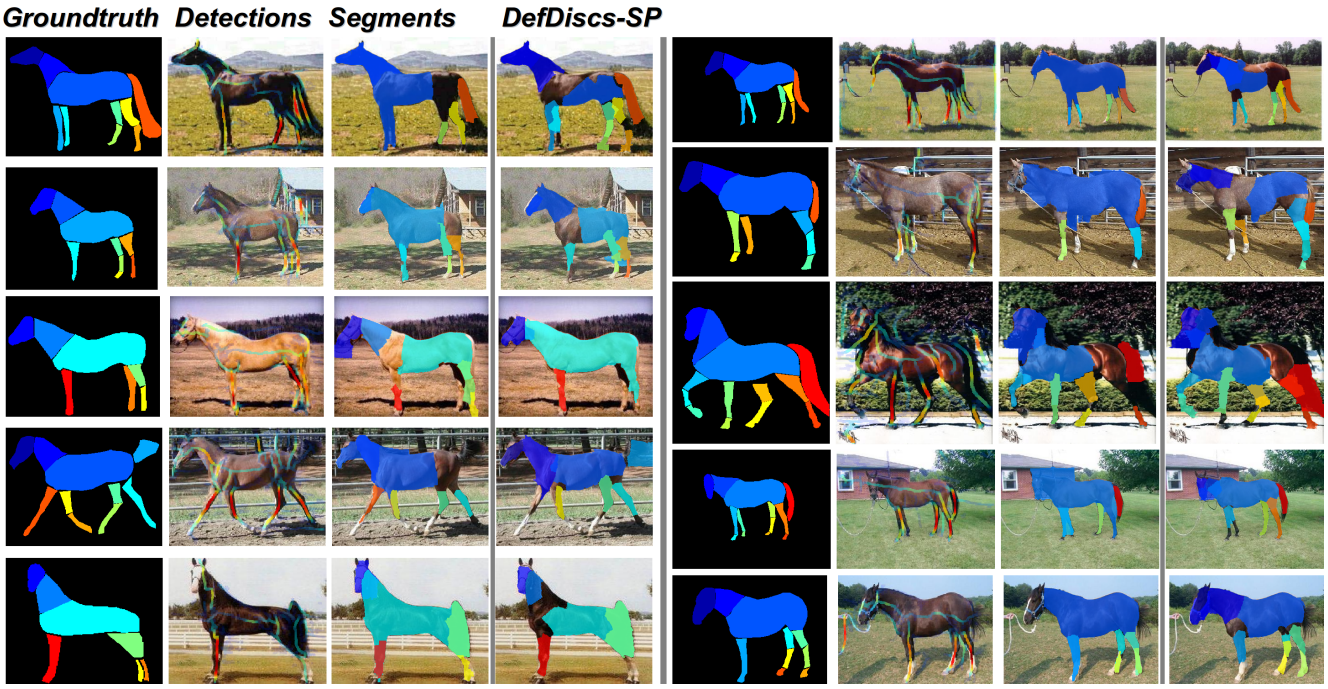

Figure 4. Additional results for WHD. 5 results per panel (L-R): Groundtruth symmetrical regions, `SRFSym` detections, `SymSegGC` segmentations, `DefDiscs-SP` segmentations. We show here segmentations that have significant overlap with the groundtruth, and are color coded accordingly.

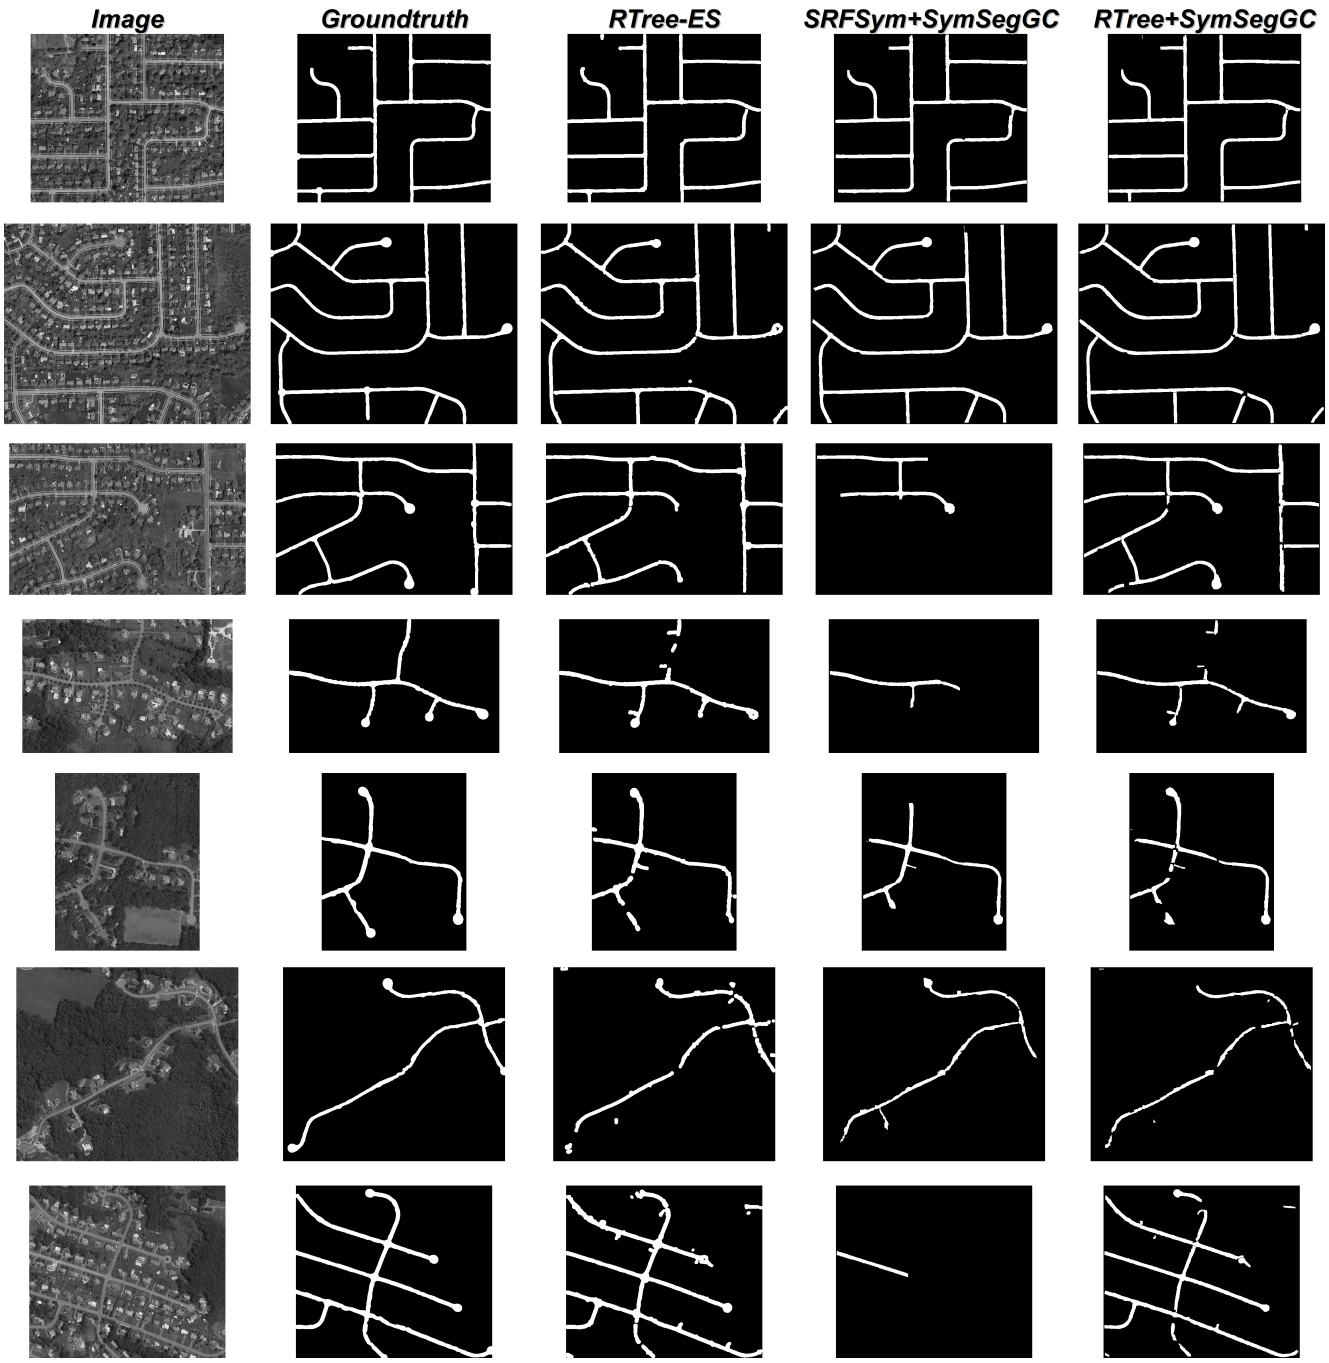

Figure 5. Additional results for NY-roads. 7 test images (L-R): Input image, groundtruth segmentations, `RTree-ES` segmentations, `SRFSym+SymSegGC` segmentations, `RTree+SymSegGC` segmentations. We show results at 0.5 recall.