Detection and Segmentation of 2D Curved Reflection Symmetric Structures

Code and data: http://www.umiacs.umd.edu/~cteo/SymmetrySegmentation

Ching L. Teo, Cornelia Fermüller and Yiannis Aloimonos – Computer Vision Lab and UMIACS, University of Maryland, College Park, MD

Abstract

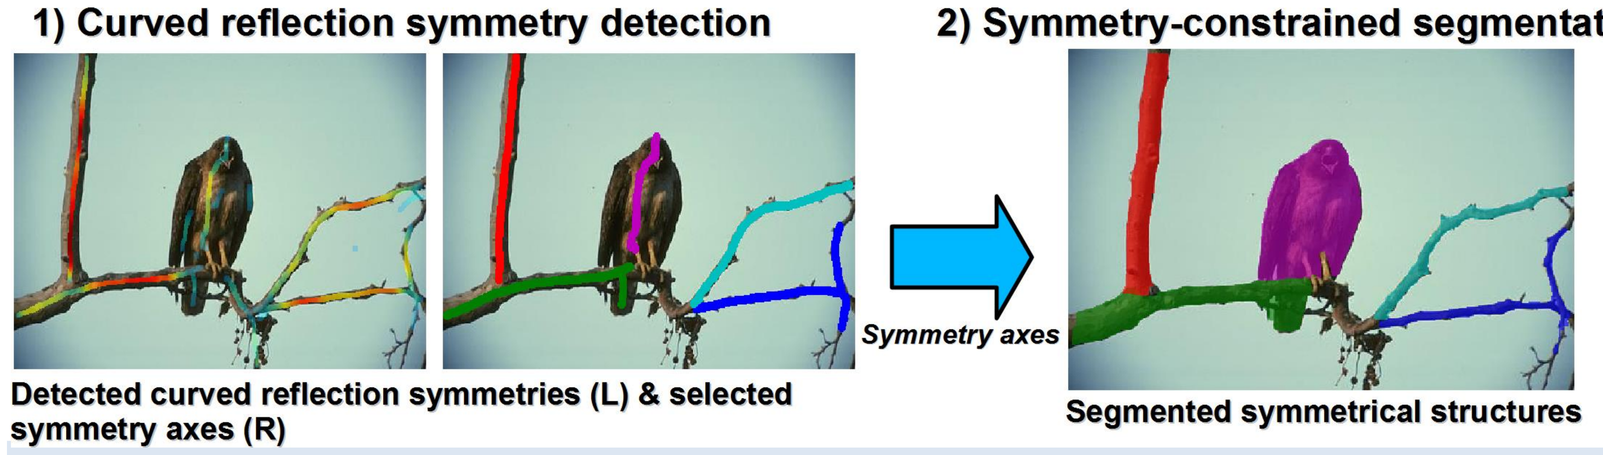

Symmetry, as one of the key tenets of Gestalt theory, is ubiquitous in natural images. In this work, we focus on 1) **detecting curved reflection symmetries** commonly found in articulated objects and 2) **segmenting symmetric regions** that support these symmetries. We do this in two steps: I) Curved symmetries are detected via a Structured Random Forest[1] (SRF) trained over local cues and II) using them as priors in a 5-way Markov Random Field (MRF) representation of the image edges to produce **symmetry-constrained** segments via graph-cuts. Consequently, our method handles 1) multiple branched symmetries and produces 2) quasi-symmetric regions in one single optimization step.

What are Curved Reflection Symmetries?

Curved reflection symmetry, also known as *ridges, ribbons or centerlines* is a generalization of reflection symmetry for articulated structures and is closely related to its *medial axis*.

Curved Symmetries: Local Cues

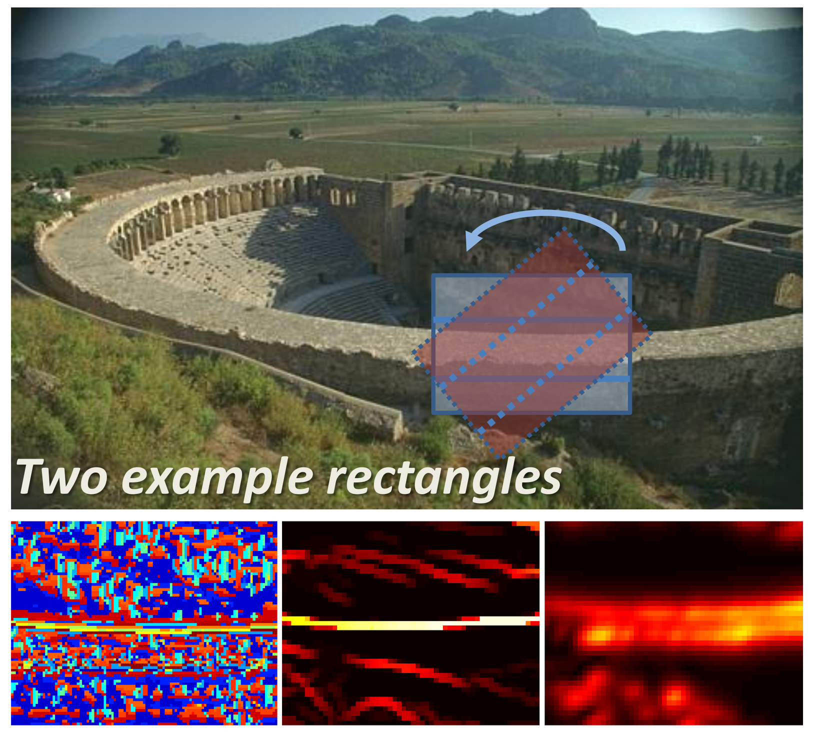

Textons *Spectral* *Gabor*

We adopt the fast integral image approach of [4] by comparing multi-scaled histograms of local cues: *L*a*b**, *textons*[3], *spectral*[4] and *Gabor* edges. These cues capture **complimentary** information – Gabors are useful when textureless regions are encountered. Histograms of these cues along oriented rectangles using the EMD-L1 distance[5] are used as input features to the SRF-based classifier.

Fast Curved Symmetry Detection via SRF

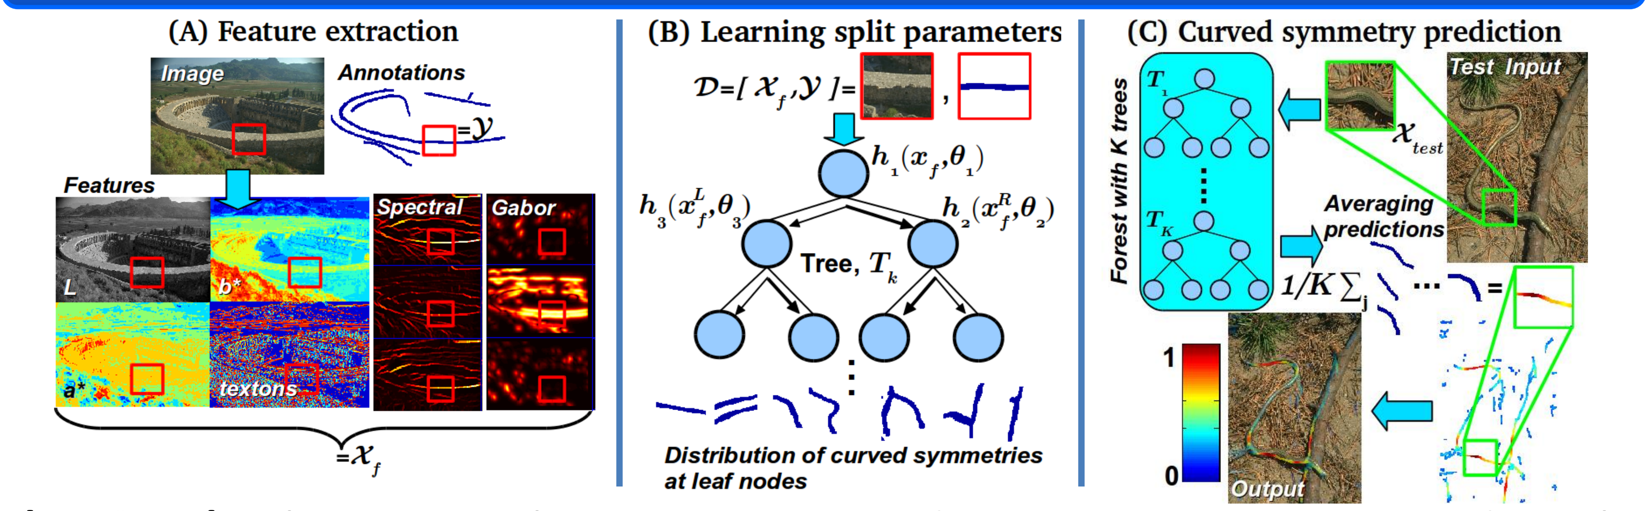

(A & B) The SRF takes in *structured annotations* – ground-truth centerlines ($\mathcal{Y}$) with input patch-based features ($\mathcal{X}_f$), and learns the optimal split parameters, $\boldsymbol{\theta}$, at each split node and a distribution of centerlines at the leaf nodes. **(C)** Given a test image, each decision tree will produce a structured prediction of the centerline given the same input features, and we average the predictions over all trees.

Symmetry-Guided Segmentation via Graph-Cuts

Given a set of axes $\boldsymbol{A}$, we construct a *5-way MRF* with an additional **cross-symmetry** connection between symmetric neighbors (p, p'):

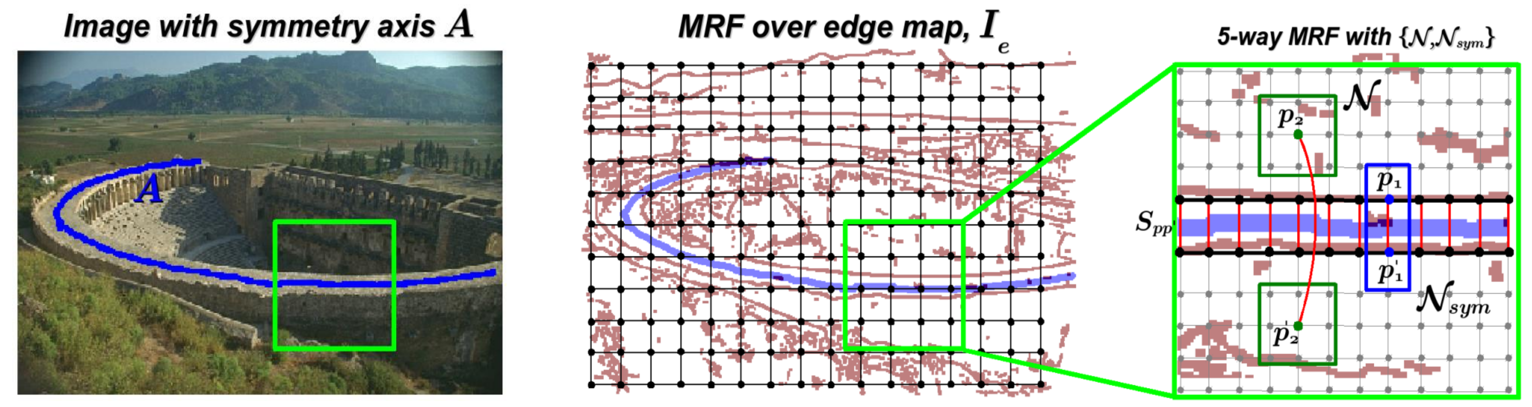

The final segmentation is obtained by minimizing the energy function:

$$E(f) = \underbrace{\sum_{p \in \mathcal{P}} U_p(f_p) + \sum_{(p,q) \in \mathcal{N}} V_{pq}(f_p, f_q)}_{\text{Unary + Pairwise (edges)}} + \underbrace{\sum_{(p,p') \in \mathcal{N}_{sym}} S_{pp'}(f_p, f_{p'}) + \sum_{(p,q) \in \mathcal{N}} B_{pq}(f_p, f_q)}_{\text{Cross-symmetry + “Ballooning”}}$$

Two **symmetry measures** between (p,p') are used: 1) the SRF-based predictions $\nu_{pp'}$ and 2) their Euclidean distance to the symmetry axis, $\mathcal{D}_{\boldsymbol{A}}$. We embed these measures into the cross-symmetry term, $S_{pp'}$.

Cross-symmetry Pairwise Term

$$S_{pp'}(f_p, f_{p'}) = \begin{cases} es_{pp'}, & \text{if } f_p \neq f_{p'} \\ \beta, & \text{otherwise} \end{cases}$$

$$es_{pp'} = 1 + \beta - \frac{1}{Z}\log(1 + \underbrace{\|D_{\mathcal{A}}(p) - D_{\mathcal{A}}(p')\| + \nu_{pp'}}_{\text{Symmetry strength + SRF score}})$$

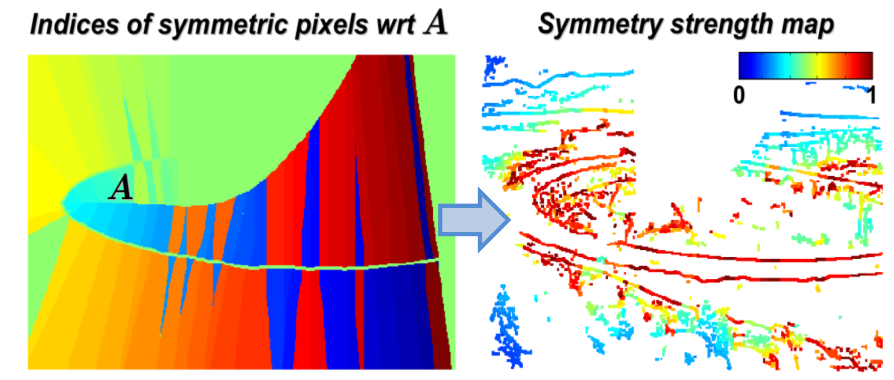

Large penalty for symmetric pixels with different labels → enforces symmetry in final segmentation.

“Ballooning” Expansion Term

Noise and internal edges (due to textures) often results in a degenerate segmentation close to the symmetry axis. Similar to [6], we encourage the segmentation to expand along *opposing directions* perpendicular to the symmetry axis.

$$B_{pq}(f_p, f_q) = \begin{cases} 0, & \text{if } f_p = f_q \\ \infty, & \text{if } f_p = 1 \text{ and } f_q = 0 \\ \rho_b, & \text{if } f_p = 0 \text{ and } f_q = 1 \end{cases}$$

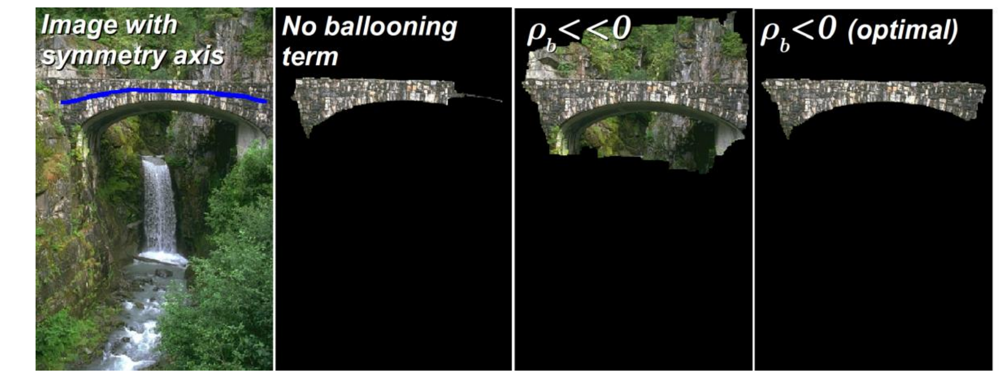

Submodularity of pairwise terms: V_{pq}, B_{pq} are submodular by construction. Since $es_{pp'} \geq \beta$ for all β, $E(f)$ can be minimized exactly via graph-cuts[7].

Experimental Evaluations

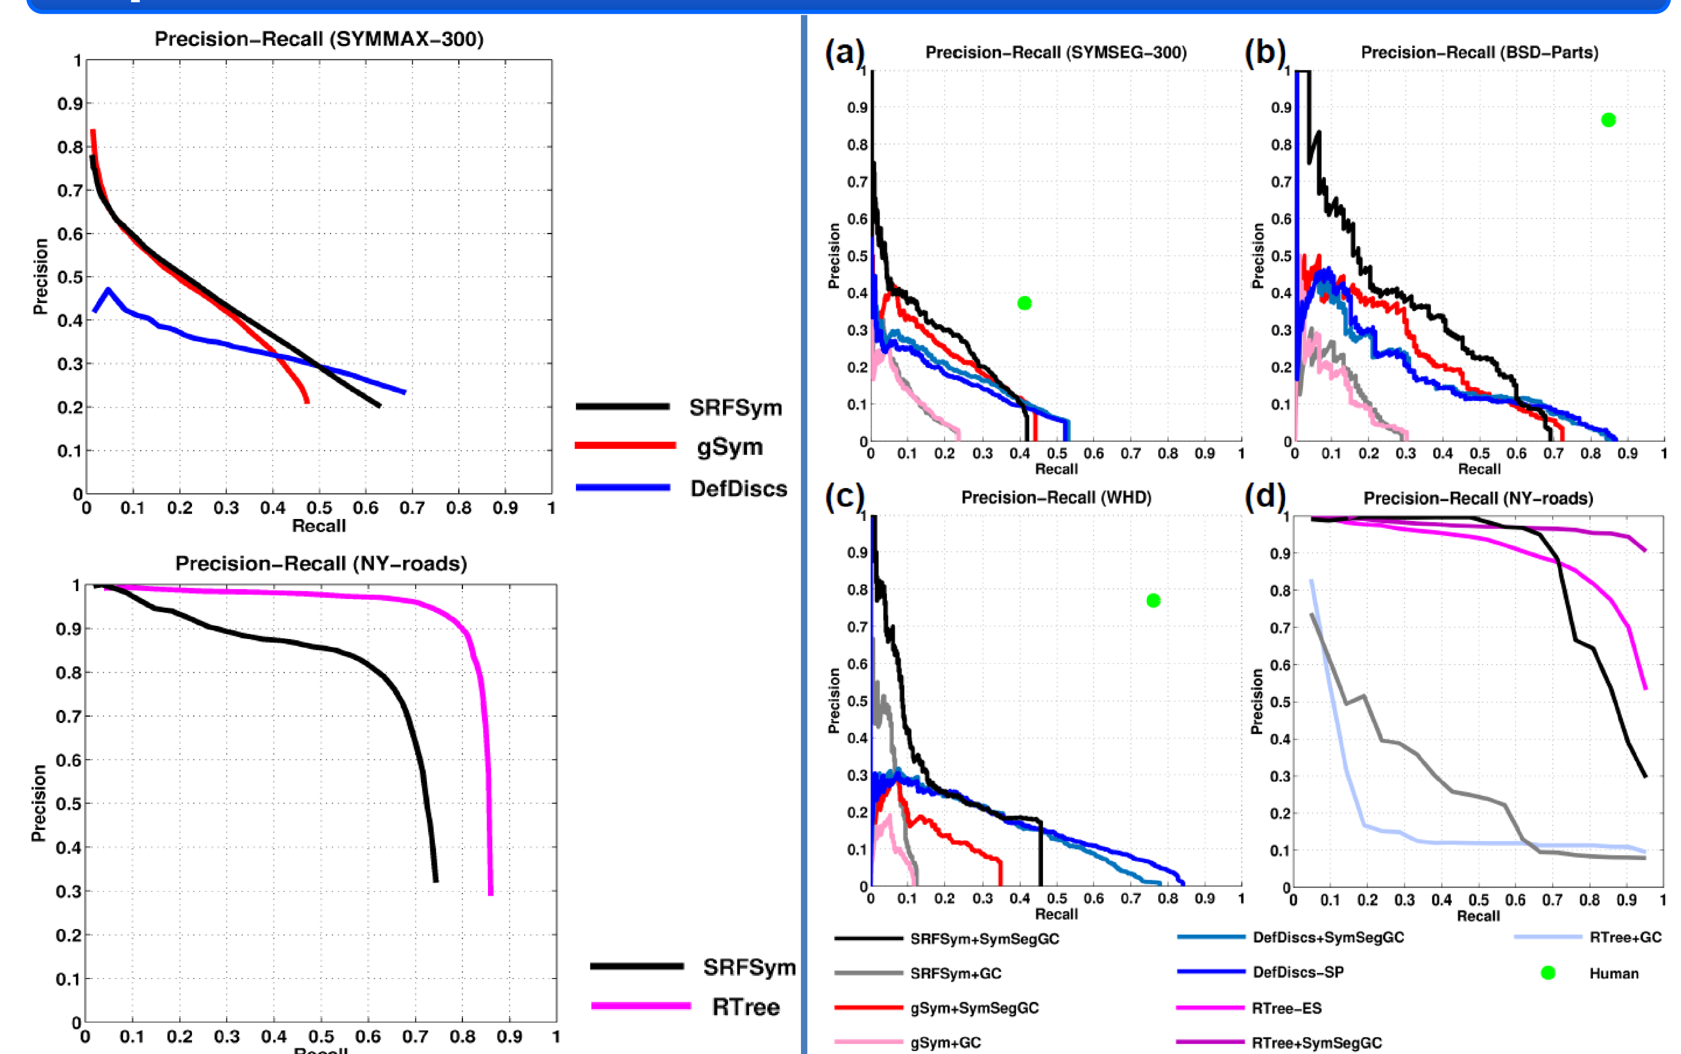

Curved Symmetry Detection **Symmetry-Constrained Segmentation**

Evaluation on two tasks: a) curved symmetry detection and b) segmentation accuracy over four publicly available datasets[4,8,9]. Our approach produces **competitive** curved symmetry predictions and **superior** symmetrical segments at a **fraction of the time** needed for other approaches.

Acknowledgements

This research was funded in part by the support of the European Union under the Cognitive Systems program (project POETICON++), the National Science Foundation under INSPIRE grant SMA 1248056 and DARPA Grant (W911NF-14-1-0384) under the project: Shared Perception, Cognition and Reasoning for Autonomy.

References

[1] P. Dollar and C. L. Zitnick. “Fast edge detection using structured forests”. *PAMI*, 37(8):1558–1570, 2015.
[2] Y. Liu, H. Hel-Or, C. S. Kaplan, and L. Van Gool. “Computational symmetry in computer vision and computer graphics”. *Foundations and Trends in Computer Graphics and Vision*, 5(12):1–195, 2009.
[3] D. R. Martin, C. C. Fowlkes, and J. Malik. “Learning to detect natural image boundaries using local brightness, color, and texture cues”. *PAMI*, 26(5):530–549, 2004
[4] S. Tsogkas and I. Kokkinos. “Learning-based symmetry detection in natural images”. In *ECCV*, pages 41–54. 2012.
[5] H. Ling and K. Okada. “An efficient earth mover’s distance algorithm for robust histogram comparison”. *PAMI*, 29(5):840–853, 2007.
[6] O. Veksler. “Star shape prior for graph-cut image segmentation”. In *ECCV*, pages 454–467. 2008.
[7] V. Kolmogorov and R. Zabih. “What energy functions can be minimized via graph cuts?” *PAMI*, 26(2):147–159, 2004.
[8] T. S. H. Lee, S. Fidler, and S. Dickinson. “Detecting curved symmetric parts using a deformable disc model”. In *ICCV*, pages 1753–1760, 2013.
[9] A. Sironi, V. Lepetit, and P. Fua. “Multiscale centerline detection by learning a scale-space distance transform”. In *CVPR*, pages 2697–2704, 2014

Scan for more Info: