

# Simultaneous Segmentation and Figure/Ground Organization using Angular Embedding

Michael Maire

California Institute of Technology - Pasadena, CA 91125  
mmaire@caltech.edu

**Abstract.** Image segmentation and figure/ground organization are fundamental steps in visual perception. This paper introduces an algorithm that couples these tasks together in a single grouping framework driven by low-level image cues. By encoding both affinity and ordering preferences in a common representation and solving an Angular Embedding problem, we allow segmentation cues to influence figure/ground assignment and figure/ground cues to influence segmentation. Results are comparable to state-of-the-art automatic image segmentation systems, while additionally providing a global figure/ground ordering on regions.

## 1 Introduction

Segmentation, the task of partitioning an image into homogeneous regions, and figure/ground organization, the task of assigning ownership of a contour to one of the two regions it separates, are both active and open problems in computer vision. Historically, more attention has been paid to segmentation, though some important studies of figure/ground exist, focusing on contour and junction structure [13, 11, 25, 32] or specific cues [10] such as convexity [21] or lower-region [29]. Recent work has revived interest on figure/ground discrimination [24, 16] and the related problem of depth ordering [15, 26].

Previous work starts from the assumption that figure/ground organization occurs after contours [24] or regions [16] have been obtained and designs algorithms that require an image segmentation as input. Hoiem *et al.* [15] fix an initial oversegmentation and iterate region-merging and depth estimation steps. It is not yet known where figure/ground discrimination occurs in biological visual systems [22], with, as noted by Ren *et al.* [24], some evidence for early availability of a contour ownership signal [34].

Most automatic image segmentation algorithms ignore figure/ground organization, producing a two-dimensional partition of the image with no notion of figure or depth ordering [28, 6, 30, 8, 27, 1, 23, 3]. Other work treats depth recovery itself as an end goal, exploiting segmentation along with scene geometry (*e.g.* estimated horizon location) or object knowledge to help build a three-dimensional rendering of the image. Examples include the photo pop-up work of Hoiem *et al.* [14] and the scene reconstruction work of Gould *et al.* [12].

This paper takes a different approach, attempting to bring figure/ground cues into perceptual processing as early as possible. We want to build a generic

segmentation and figure/ground reasoning stage with the goal of enriching the image representation available to tasks such as object recognition.

We accomplish this by extending a leading image segmentation method based on spectral partitioning into an algorithm that recovers figure/ground organization as well. The system we extend is that of Arbeláez *et al.* [3], which currently provides the best performance of all automatic segmentation algorithms across a range of benchmarks on the Berkeley Segmentation Dataset (BSDS) [19, 18]. Our key insight is to replace their core grouping machinery, based on Normalized Cuts [28] and described in Maire *et al.* [17], with the more general Angular Embedding of Yu [31]. Angular Embedding allows us to represent both segmentation and figure/ground relations and solve for both at once.

To our knowledge, no previous work recovers segmentation and figure/ground for natural images in a single step. Yu and Shi [33] attempt to use pairwise repulsion cues to fuse figure/ground with segmentation in spectral graph theory. However, they show only one example on T-junctions. A core component of our solution, Angular Embedding, was only recently introduced [31] and we believe our work is the first application of this technique to non-synthetic images.

The most closely related work to ours is that of Ren *et al.* [24] and Leichter and Lindenbaum [16], both of which focus on solving an easier problem than the one suggested here. Leichter and Lindenbaum take the human-drawn ground-truth segmentations [19] and human figure/ground annotations [10] of the BSDS images and learn a conditional random field (CRF) for assigning boundary ownership. They use curve and junction potentials, exploiting convexity, lower-region, fold/cut, and parallelism cues. Their impressive results of 82.8% correct figure/ground assignments (chance being 50%) are only obtained when *testing on human-drawn ground-truth segmentations*. Testing on automatically generated curves, they obtain only 69.1% accuracy, similar to the 68.9% accuracy reported by Ren *et al.* [24] on automatically generated contours. In contrast to our integrated approach, both works cast figure/ground assignment as a step to be run after first solving for a segmentation.

The notion of figure/ground used in this work is that used by Ren *et al.* [24]. Namely, a figural region is defined according to human perception. We simply attempt to replicate human behavior by training on human-annotated data. As previously pointed out [20, 16], this means that figure/ground ordering does not necessarily correspond to depth or occlusion ordering. For example, humans may indicate strong figure/ground percepts due to markings on flat surfaces.

Our choice to follow this convention for the meaning of figural regions is consistent with the goal of targeting our output for use in perceptual tasks such as recognition rather than geometric scene reconstruction. It is also partially motivated by convenience, as it allows us to train our figure/ground classifier on the same dataset, the BSDS, as our segmentation algorithm, due to the availability of pre-existing annotations [19, 10]. Consequently, our work is not directly comparable to that for which depth ordering or three-dimensional reconstruction is the ultimate goal [15, 26]. However, since our algorithm for the combined segmentation and figure/ground problem is agnostic to the source of the local fig-

ure/ground cues, it is conceivable that future work could re-purpose our system to solve a depth ordering problem.

Section 2 describes our new grouping framework for simultaneous segmentation and figure/ground assignment. It is compatible with any appropriate sources of pairwise similarity and ordering cues. Section 3 details our particular choice for the local ordering cues. Section 4 presents both qualitative and quantitative results for fully automatic segmentation and figure/ground organization. Our system compares favorably to others on the segmentation task, while producing a global figural ranking of regions at minimal additional computational cost.

## 2 Adding Ordering to Segmentation

Figure 1 outlines our algorithm for simultaneous segmentation and figure/ground organization. We extend previous work on segmentation alone [17, 3] to incorporate figure/ground information through the use of Angular Embedding [31] as a globalization procedure. Removing the vertical pathway for figure/ground information shown on the right side of Figure 1 and replacing Angular Embedding by Normalized Cuts, one recovers the segmentation-only pipeline of Arbeláez *et al.* [3]. To make this paper as self-contained as possible, we briefly review the core portions of the relevant previous work, before describing how to bring in figure/ground cues in the form of pairwise ordering preferences.

### 2.1 Spectral Partitioning

Spectral clustering, and specifically Normalized Cuts [28], have long been popular techniques for image segmentation. Recently, [17] achieve excellent results by using Normalized Cuts in a “soft” manner as a globalization stage for contour detection. The approach taken is to define a sparse affinity matrix connecting nearby pixels  $p$  and  $q$  with weight determined by the *intervening contour* [9] cue:

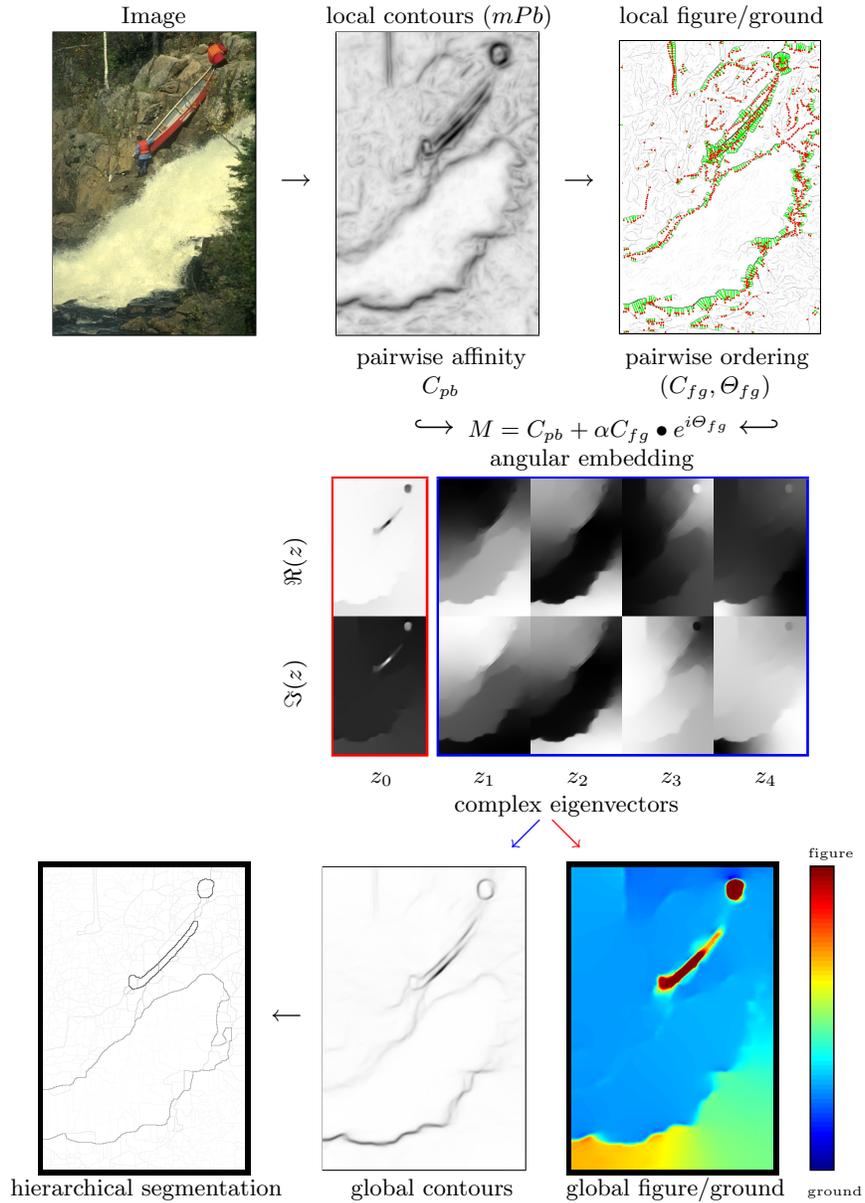
$$W(p, q) = \exp \left( - \max_{(x,y) \in \overline{pq}} \{mPb(x,y)\} / \rho \right) \quad (1)$$

where  $\overline{pq}$  is the line segment connecting  $p$  and  $q$ ,  $\rho$  is a constant, and  $mPb$  stands for multiscale *probability of boundary* [18, 17] and measures the probability that the pixel at location  $(x, y)$  lies on a boundary contour. A classifier trained using local brightness, color, and texture cues predicts  $mPb$  at each image location.

To obtain global contour strength from these local measurements, one forms matrix  $D$  whose diagonal contains the row-sums of  $W$  and solves for the generalized eigenvectors  $\{v_0, v_1, \dots, v_n\}$  of the system:

$$(D - W)v = \lambda Dv \quad (2)$$

corresponding to the  $n + 1$  smallest eigenvalues  $0 = \lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_n$ . Associating with each pixel  $p$  the length  $n$  descriptor containing the  $p^{\text{th}}$  entry from each eigenvector creates an embedding in  $\mathbb{R}^n$ . Equivalently,  $\{v_1, \dots, v_n\}$  can be



**Fig. 1. Segmentation and figure/ground organization.** From the image (*top left*) we compute the probability of boundary ( $pb$ ) [18] using the multiscale detector ( $mPb$ ) of [17] (*top middle*). Nonmax-suppressed  $mPb$  contours are fed to a local shape-based figure/ground classifier (*top right*), whose output is shown by green vectors with red tips drawn towards the predicted figural side. The  $mPb$  signal defines a pairwise affinity between neighboring pixels via intervening contour [9]. The figure/ground classifier defines a longer-range pairwise ordering. A generalized affinity matrix  $M$  captures both sources of information. Solving an Angular Embedding [31] problem yields complex eigenvectors (*middle*) which encode both segmentation (*bottom left*) and global figure/ground ordering (*bottom right*). Red indicates more figural regions.

viewed as a stack of  $n$  images for which the segmentation problem is now easy. Convolution with Gaussian directional derivative filters produces a robust measure of contour strength. Applying tools from image morphology then permits construction of a hierarchical segmentation from these high-quality contours [3].

For an intuition behind this machinery, note that the first eigenvector  $v_1$  is the exact global minimizer of the following error measure [28]:

$$\inf_{v^T D 1=0} \frac{\sum_p \sum_q W(p, q) (v(p) - v(q))^2}{\sum_p D(p) (v(p))^2} \quad (3)$$

The weight on the squared difference forces the eigenvector to take similar values for pixels with high affinity.

## 2.2 Angular Embedding

The spectral partitioning algorithm of the previous section produces real-valued eigenvectors. Angular Embedding [31] is an alternative technique that produces complex-valued eigenvectors. Our problem is no longer defined by the symmetric real-valued matrix  $W$ , but instead by a pair of real-valued matrices  $(C, \Theta)$ , where  $C$  is a symmetric *confidence* matrix analogous to  $W$ , and  $\Theta$  is a skew-symmetric *ordering* matrix. The goal is to produce an embedding into the unit circle in the complex plane such that sorting the resulting points by their angle respects the pairwise local ordering constraints defined by  $\Theta$ . Confidence matrix  $C$  encodes the relative importance of each constraint.

Specifically, let  $z(p) \in \mathbb{C}$  denote the embedding of  $p$ . We minimize error:

$$\varepsilon = \sum_p D(p) \cdot |z(p) - \tilde{z}(p)|^2 \quad (4)$$

where  $D$  is again a diagonal degree matrix with:

$$D(p) = \frac{\sum_q C(p, q)}{\sum_{p, q} C(p, q)} \quad (5)$$

and  $\tilde{z}(p)$  is the position of  $p$  estimated from its neighbors through a rotation by their relative ordering:

$$\tilde{z}(p) = \sum_q \tilde{C}(p, q) \cdot e^{i\Theta(p, q)} \cdot z(q) \quad (6)$$

$$\tilde{C}(p, q) = \frac{C(p, q)}{\sum_q C(p, q)} \quad (7)$$

$|z - \tilde{z}|$  is an appropriate error measure as  $z$  and  $\tilde{z}$  coincide if and only if the embedding perfectly fulfills all local orderings with positive confidence [31].

Rewriting the above in matrix form requires one to minimize:

$$\varepsilon = z^* W z \quad (8)$$

subject to  $z = e^{i\theta}$  for a real-valued vector  $\theta$  where:

$$W = (I - M)^* D (I - M) \quad (9)$$

$$M = \text{Diag}(C1)^{-1} C \bullet e^{i\Theta} \quad (10)$$

$$D = \text{Diag}(C1 \cdot (1^* C1)^{-1}) \quad (11)$$

and  $*$  denotes complex conjugate transpose,  $\bullet$  is the matrix Hadamard product,  $I$  is the identity matrix,  $1$  is a column vector of ones,  $\text{Diag}(\cdot)$  is a matrix with its vector argument on the main diagonal,  $i = \sqrt{-1}$  and exponentiation acts element-wise. Relaxing the constraint that  $z$  lie on the unit circle to  $z^* D z = 1$  yields the solution as the angle of the first eigenvector,  $\angle z_0$ , of the generalized eigenproblem specified by  $(W, D)$ . Unlike (2), for nontrivial  $\Theta$ , we have  $\lambda_0 \neq 0$  and all of the eigenvectors, including  $z_0$ , are meaningful.

### 2.3 Short-Range Attraction, Long-Range Ordering

We use the additional expressive freedom of Angular Embedding to encode both pairwise segmentation cues and pairwise figure/ground cues in the common representation defined by  $(C, \Theta)$ . Let us now write the affinity matrix defined by intervening contour (1) as  $C_{pb}(p, q)$ . It uses the probability of boundary ( $pb$ ) cue to place a confidence on the event that pixels  $p$  and  $q$  lie in the same region. This cue yields no information on relative figural ordering, so we set  $\Theta_{pb}(p, q) = 0 \forall p, q$ .

Suppose we also have a classifier  $f(x, y) \rightarrow [-1, 1]$  that, at an edge pixel  $(x, y)$  lying on a contour obtained by nonmax-suppression [5] of  $mPb$ , predicts which side of the edge is figural. Let  $p$  and  $q$  be the pixels located a fixed distance  $r$  from  $(x, y)$ , on opposite sides (left and right, respectively) of the edge, in the direction orthogonal to the local edge orientation, as shown in Figure 2. Define:

$$C_{fg}(p, q) = C_{fg}(q, p) = |f(x, y)| \cdot mPb(x, y) \quad (12)$$

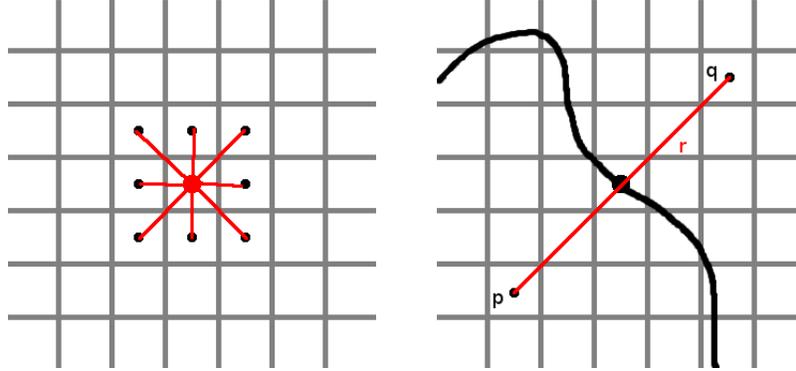
$$\Theta_{fg}(p, q) = -\Theta_{fg}(q, p) = \text{sign}(f(x, y)) \cdot \phi \quad (13)$$

where  $\phi$  represents a constant angular separation. These equations state that  $p$  and  $q$  should be embedded at angular separation  $\phi$  with confidence that increases with figure/ground classifier confidence and edge strength.  $\phi$  must be chosen sufficiently small such that the number of figure/ground layers in the image does not exceed  $\frac{\pi}{\phi}$ . We set  $\phi = \frac{\pi}{8}$  in experiments.

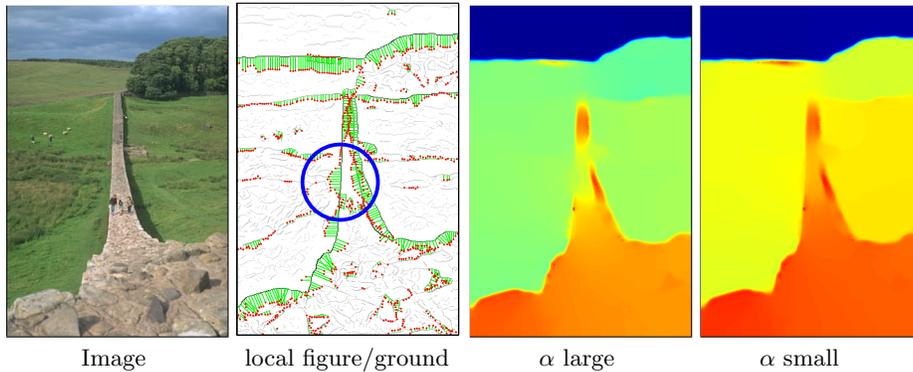
By choosing  $r$  greater than the radius used for local intervening contour affinities,  $(C_{pb}, \Theta_{pb})$  and  $(C_{fg}, \Theta_{fg})$  have no overlapping nonzero entries. Writing  $C = C_{pb} + \alpha C_{fg}$  and  $\Theta = \Theta_{pb} + \Theta_{fg}$ , with  $\alpha$  weighting the relative importance of the two signals, a measurement matrix  $M$  captures all information (excluding the normalization term involving  $C$ ):

$$M = C \bullet e^{i\Theta} = C_{pb} \bullet e^{i\Theta_{pb}} + \alpha C_{fg} \bullet e^{i\Theta_{fg}} = C_{pb} + \alpha C_{fg} \bullet e^{i\Theta_{fg}} \quad (14)$$

The short-range connections  $(C_{pb}, \Theta_{pb})$  encode the prior that there is no figure/ground difference between neighboring pixels, but the confidence in this



**Fig. 2. Pairwise attraction and ordering.** *Left:* We connect each pixel to its 8 immediate neighbors with affinity depending on the computed edge strength ( $pb$ ) between them (the intervening contour [9]). A sparse matrix  $C_{pb}$  encodes these preferences. *Right:* The figure/ground classifier runs on a sampled set of nonmax-suppressed edge pixels. In each case, it induces a connection between the two pixels  $p$  and  $q$  located a fixed distance  $r$  from the edge point, in the direction orthogonal to the edge orientation. The predicted figural side defines a relative ordering  $\Theta_{fg}(p, q)$ , with an associated confidence  $C_{fg}(p, q)$ . Measurement matrix  $M = C_{pb} + \alpha C_{fg} \bullet e^{i\Theta_{fg}}$  (up to a normalization factor) encodes both types of information, where  $\bullet$  denotes element-wise product.



**Fig. 3. Competing segmentation and figure/ground cues.** Where segmentation and local figure/ground predictions disagree (*blue circle*), the relative weighting,  $\alpha$ , of the cues determines which dominates. For large  $\alpha$ , locally incorrect figure/ground classification (*middle left*) overrules the tendency towards coherent ordering within regions, resulting in incorrect figure/ground globalization (*middle right*). For smaller  $\alpha$ , adherence to strong segment boundaries corrects local figure/ground errors (*far right*). Setting  $\alpha = 0$  results in recovery of segmentation only and ignores figure/ground ordering (*not shown*).

prior decreases in the vicinity of a strong edge. The longer-range connections  $(C_{fg}, \Theta_{fg})$  encode relative figure/ground ordering between more distant pixels. Attempting to satisfy both constraints yields an embedding which can violate the uniform prior near boundaries (where its confidence is low), in order to break apart figure and ground regions. Conversely, figure/ground differences are suppressed within a segment. Figure 3 demonstrates this type of competition.

## 2.4 Eigenvector Interpretation

Recall from Section 2.2 that the angle of the leading complex-valued eigenvector,  $\angle z_0$ , assigns each pixel a global figural ordering when solving the Angular Embedding problem specified by (14). We are left with the question of how to extract a segmentation. Note that in the absence of figure/ground cues,  $C_{fg} = 0$  and  $\Theta_{fg} = 0$ ,  $M$  is real-valued, and we find  $\angle z_0 = 0$ . Looking at the first  $n+1$  eigenvectors  $\{z_0, z_1, \dots, z_n\}$  and their corresponding eigenvalues,  $\lambda_0 \leq \lambda_1 \leq \dots \leq \lambda_n$ , we find a similar situation as for Normalized Cuts (2). In particular,  $\lambda_0 = 0$ , but the remaining eigenvectors provide an embedding in which segmentation is easy. For the general case with figure/ground cues, the now complex-valued eigenvectors still provide such an embedding (though with  $\lambda_0 \neq 0$  and  $z_0$  nontrivial).

We therefore extend the idea of extracting contours by computing gradients on the stack of eigenvector images [17] to the complex-valued case. As Figure 4 shows, we compute the “spectral” contour signal:

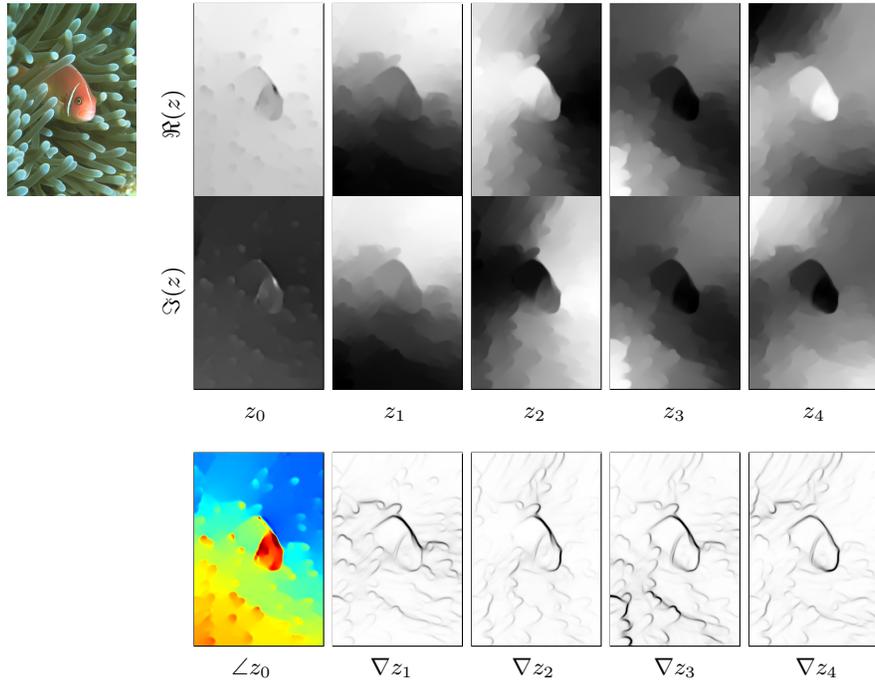
$$sPb(x, y, \theta) = \sum_{k=1}^n \frac{1}{\sqrt{\lambda_k}} \cdot \left( [\nabla_{\theta} \Re\{z_k(x, y)\}]^2 + [\nabla_{\theta} \Im\{z_k(x, y)\}]^2 \right)^{\frac{1}{2}} \quad (15)$$

Following the procedure of Arbeláez *et al.* [3], we create a weighted combination of  $mPb$  and  $sPb$  and apply their Oriented Watershed Transform - Ultrametric Contour Map (OWT-UCM) algorithm to construct a hierarchical image segmentation. Averaging  $\angle z_0$  over the resulting segments translates our figure/ground ordering on pixels into an ordering on regions.

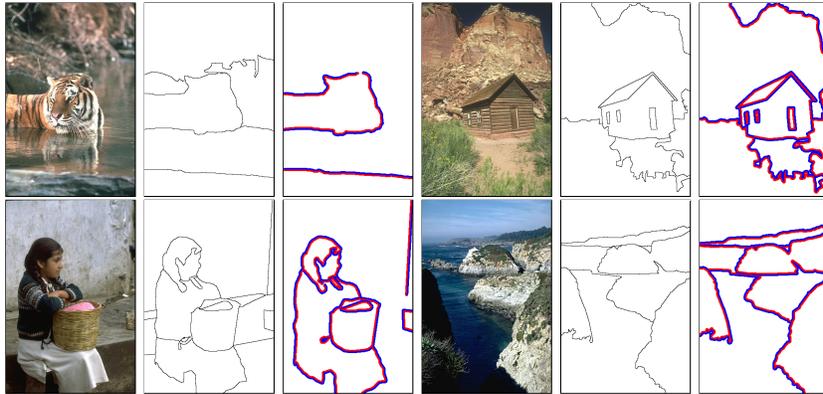
## 3 Local Figure/Ground Classifier

The presentation so far has omitted the details of the figure/ground classifier introduced in Section 2.3. As previously mentioned, this classifier could predict depth ordering or a use a perceptual notion of figuralness. We choose the latter definition for convenience, keeping in mind that the primary focus of this paper is the new globalization algorithm for coupling figure/ground organization to segmentation, and not the engineering of this local classifier.

Rather than hand-design features for the figure/ground classifier, we borrow the approach of Ren *et al.* [24] and compute Geometric Blur [4] descriptors on top of the local  $mPb$  contour signal. We rotate each descriptor according to a local orientation estimate in an imperfect attempt to build-in limited rotation invariance (our final learned figure/ground classifier is not fully rotation



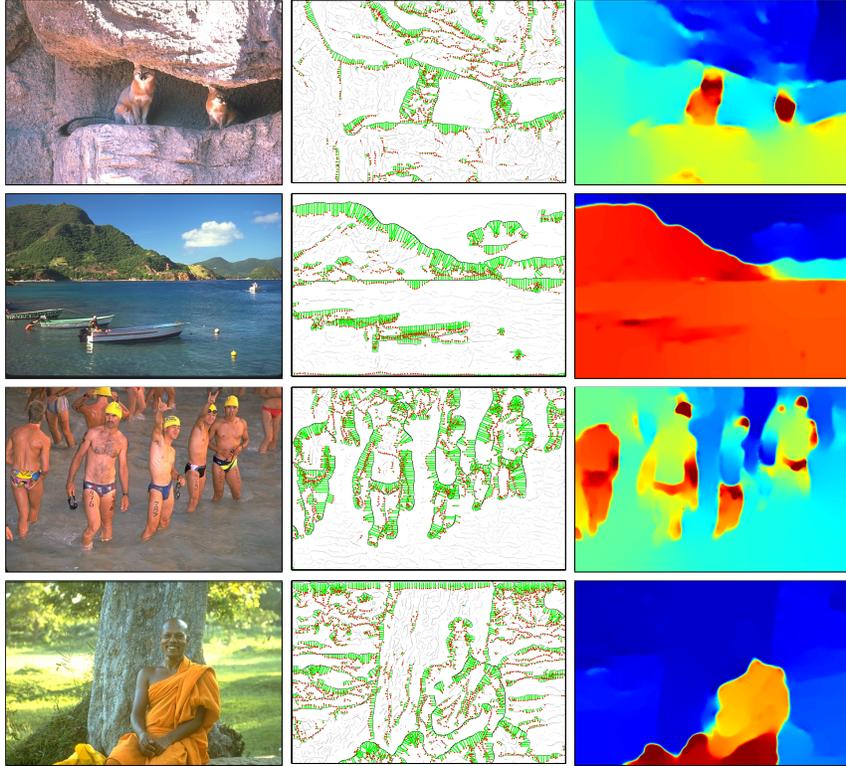
**Fig. 4. Extracting figure/ground and segmentation from eigenvectors.** *Top:* Real and imaginary components of the first five generalized eigenvectors,  $z_0, \dots, z_4$  obtained via Angular Embedding [31]. *Bottom Left:* Global figure/ground ordering is reported by  $\angle z_0$ . *Bottom Right:* Maximum oriented gradients of eigenvectors,  $\nabla z_k = \max_{\theta} \{ ([\nabla_{\theta} \Re\{z_k(x, y)\}]^2 + [\nabla_{\theta} \Im\{z_k(x, y)\}]^2)^{\frac{1}{2}} \}$ , encode a global contour signal (shown here) from which we construct a segmentation.



**Fig. 5. Berkeley segmentation dataset (BSDS) with figure/ground labeling.** *Left to Right:* Image, segment boundaries, and figure/ground annotations on a subset of those boundaries according to a human subject. Red marks the figural side. We use pre-existing segment [19] and figure/ground [10] labeling.

invariant). Clustering these descriptors using  $K$ -means (with  $K = 64$ ) yields a vocabulary of shapemes [24], which capture local contour configuration. A point of interest on a test contour is described by the vector measuring the similarity of its Geometric Blur descriptor to each of the shapemes.

We transfer human figure/ground labeling to the automatically generated nonmax-suppressed  $mPb$  contours using bipartite matching of edge pixels. We then train a logistic-regression classifier  $f$  that predicts local figure/ground assignment using the vector of shapeme similarities. This learned classifier performs at 62% accuracy, similar to the 65% accuracy reported by Ren *et al.* [24] for their local classifier. Figure 5 shows example human-annotated training data for this task and Figure 6 demonstrates local figure/ground predictions and recovered global ordering. In order to take only fairly reliable predictions into account during globalization, we sample edge locations  $(x, y)$  for which  $mPb(x, y) > \tau$  and only run the local figure/ground classifier at those locations. We set  $\tau = 0.3$ .



**Fig. 6. Local to global figure/ground.** *Left:* Image. *Middle:* Local figure/ground assignment by our shape-based classifier for the most salient  $mPb$  [17] contours. Vectors drawn from edge points indicate the predicted figural side by their red tip. Vector length corresponds to classifier confidence. *Right:* Recovered global figural ordering.

## 4 Experiments and Discussion

Figures 7 and 8 show output of our algorithm for segmentation and figure/ground ordering on images from the Berkeley segmentation dataset (BSDS). Figure 9 compares our automatically generated segmentations to those of other algorithms [8, 7, 6, 2, 3] using the standard BSDS boundary precision-recall benchmark [18]. Precision-recall curves for the other algorithms are those reported in [3]. Our segmentations are better than all except those of the leading *gPb*-owt-ucm algorithm [3], to which they are fairly close. Though our algorithm can be seen a generalization of *gPb*-owt-ucm, there are a few technical differences in our implementation that may explain the small performance gap. One such difference is that we compute affinities only between neighboring pixels, reserving long-range connections for ordering cues, rather than use a larger radius for the intervening contour computation.

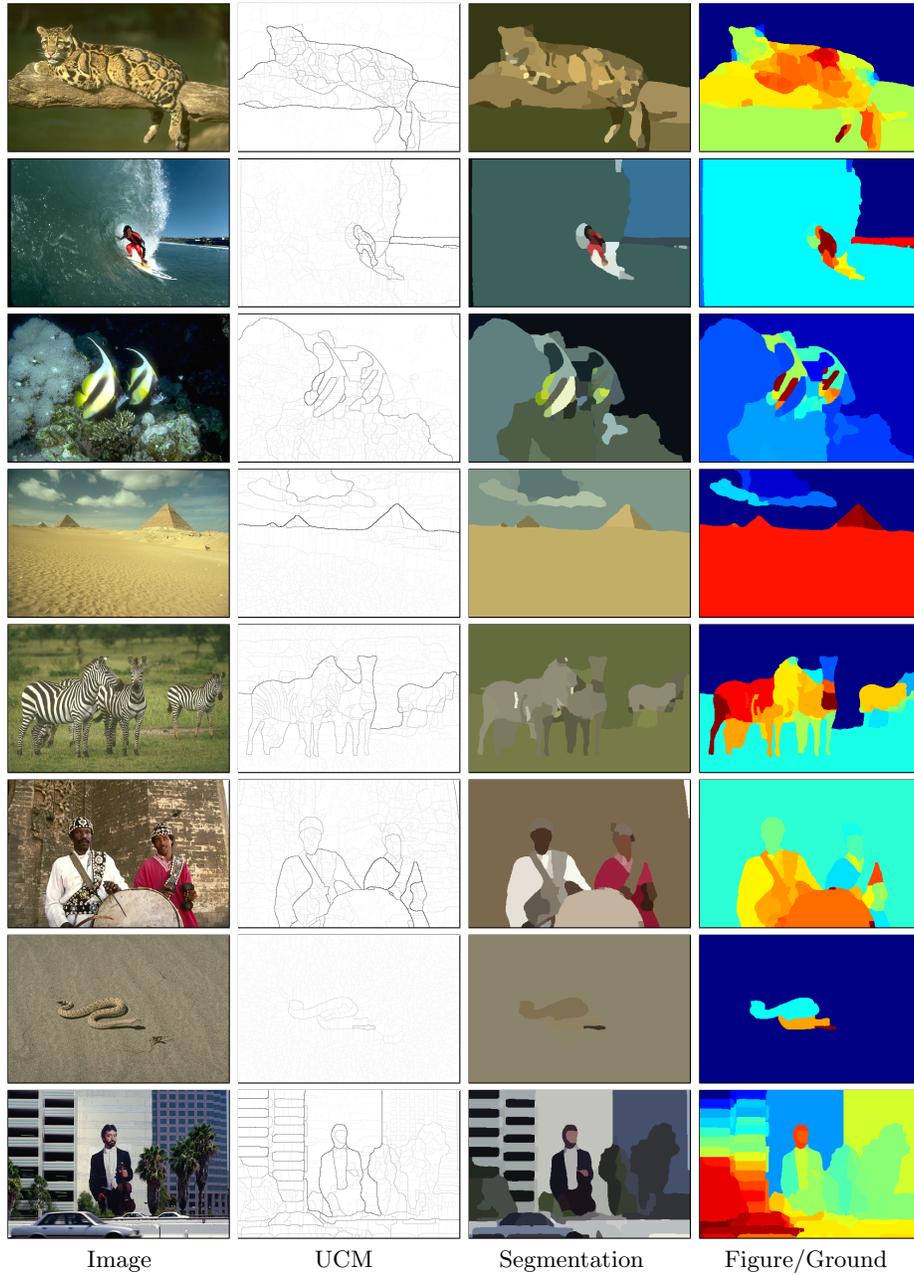
Not captured by these benchmarks is the fact that our system is the only one to solve for figure/ground. Moreover, our figure/ground output is not just a local determination of the figural side of each boundary, but is a global ranking of the regions in the segmentation. Our system offers the benefit of transforming a local figure/ground property defined on contours into a global one defined on regions. This may prove useful as a salience measure. For example, using only our bottom-up cues, the face automatically pops out as a figural region in the last example in Figure 7.

Our global figure/ground ordering comes at minimal additional computation cost over the segmentation-only approach. The local figure/ground classifier is fast to run on sampled edge points and computing eigenvectors for Angular Embedding is of the same complexity as computing them for Normalized Cuts.

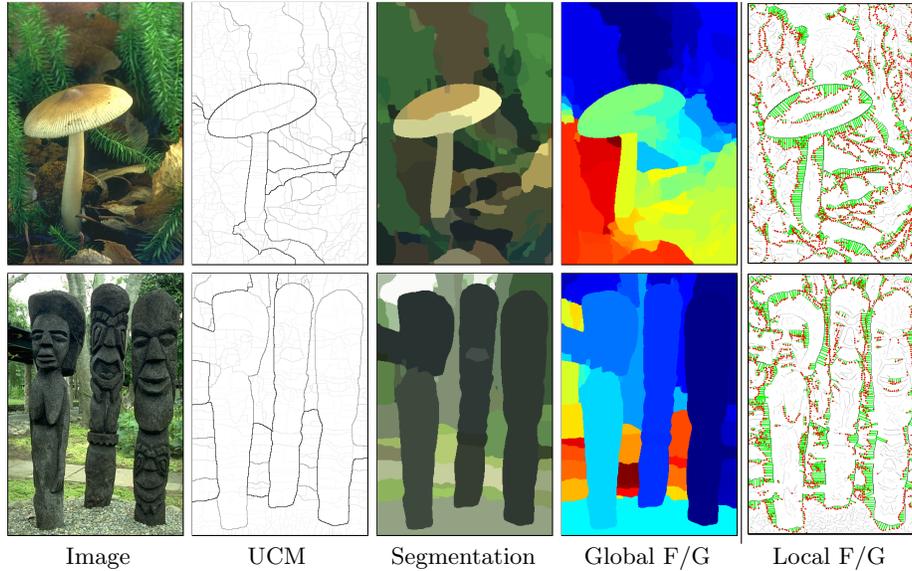
**Acknowledgments** Thanks to Stella X. Yu and Pietro Perona for helpful discussions.

## References

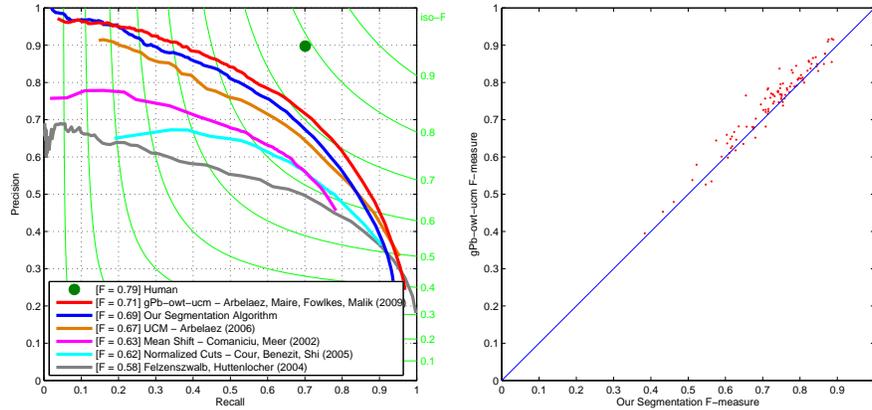
1. Alpert, S., Galun, M., Basri, R., Brandt, A.: Image segmentation by probabilistic bottom-up aggregation and cue integration. CVPR (2007)
2. Arbeláez, P.: Boundary extraction in natural images using ultrametric contour maps. POCV (2006)
3. Arbeláez, P., Maire, M., Fowlkes, C., Malik, J.: From contours to regions: An empirical evaluation. CVPR (2009)
4. Berg, A.C., Malik, J.: Geometric blur for template matching. CVPR (2001)
5. Canny, J.: A computational approach to edge detection. PAMI (1986)
6. Comaniciu, D., Meer, P.: Mean shift: A robust approach toward feature space analysis. PAMI (2002)
7. Cour, T., Benezit, F., Shi, J.: Spectral segmentation with multiscale graph decomposition. CVPR (2005)
8. Felzenszwalb, P., Huttenlocher, D.: Efficient graph-based image segmentation. IJCV (2004)



**Fig. 7. Hierarchical segmentation and figure/ground ordering results.** Our algorithm simultaneously generates a hierarchical image segmentation and a global figural ranking of regions. *Left:* Image. *Middle Left:* Hierarchical segmentation represented as an Ultrametric Contour Map (UCM) [2]. *Middle Right:* Regions at the optimal segmentation threshold displayed with their average color. *Right:* Global figure/ground ordering of the same regions. Red indicates more figural. All images shown belong to the test set.



**Fig. 8. Figure/ground failure examples.** *Left to Right:* Image, UCM, segmentation, global and local figure/ground. Globalization errors occur when the local figure/ground classifier is consistently wrong over long contours (e.g. the left side of the mushroom or the sides of the statues). Note that good segment boundaries are still recovered.



**Fig. 9. Evaluation of region boundaries on the BSDS Benchmark.** *Left:* The segmentation quality of our algorithm is close to that of the current best-performing algorithm, *gPb-owt-ucm* [3], and superior to others [8, 7, 6, 2], as benchmarked by [3]. Algorithms are evaluated in terms of precision and recall with respect to human ground-truth boundaries. The maximum F-measure ( $\frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$ ) is a summary score. Iso-F curves are shown in green. The dot indicates average human agreement. Our system is the only one that also solves for figure/ground. *Right:* Plotting per-image F-measures shows our segmentations to be competitive with those of *gPb-owt-ucm*.

9. Fowlkes, C., Martin, D., Malik, J.: Learning affinity functions for image segmentation: Combining patch-based and gradient-based approaches. *CVPR* (2003)
10. Fowlkes, C., Martin, D., Malik, J.: Local figure/ground cues are valid for natural images. *Journal of Vision* (2007)
11. Geiger, D., Kumaran, K., Parida, L.: Visual organization for figure/ground separation. *CVPR* (1996)
12. Gould, S., Fulton, R., Koller, D.: Decomposing a scene into geometric and semantically consistent regions. *ICCV* (2009)
13. Heitger, F., von der Heyd, R.: A computational model of neural contour processing: Figure-ground segregation and illusory contours. *ICCV* (1993)
14. Hoiem, D., Efros, A.A., Hebert, M.: Automatic photo pop-up. *SIGGRAPH* (2005)
15. Hoiem, D., Stein, A.N., Efros, A.A., Hebert, M.: Recovering occlusion boundaries from a single image. *ICCV* (2007)
16. Leichter, I., Lindenbaum, M.: Boundary ownership by lifting to 2.1D. *ICCV* (2009)
17. Maire, M., Arbeláez, P., Fowlkes, C., Malik, J.: Using contours to detect and localize junctions in natural images. *CVPR* (2008)
18. Martin, D., Fowlkes, C., Malik, J.: Learning to detect natural image boundaries using local brightness, color and texture cues. *PAMI* (2004)
19. Martin, D., Fowlkes, C., Tal, D., Malik, J.: A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. *ICCV* (2001)
20. Palmer, S.: *Vision Science - From Photons to Phenomenology*. MIT Press (1999)
21. Pao, H.K., Geiger, D., Rubin, N.: Measuring convexity for figure/ground separation. *ICCV* (1999)
22. Peterson, M.A., Gibson, B.S.: Must figure-ground organization precede object recognition? An assumption in peril. *Psychological Science* (1994)
23. Rao, S., Mobahi, H., Yang, A., Sastry, S., Ma, Y.: Natural image segmentation with adaptive texture and boundary encoding. *ACCV* (2009)
24. Ren, X., Fowlkes, C., Malik, J.: Figure/ground assignment in natural images. *ECCV* (2006)
25. Saund, E.: Perceptual organization of occluding contours of opaque surfaces. *CVIU Special Issue on Perceptual Organization* (1999)
26. Saxena, A., Chung, S.H., Ng, A.Y.: 3-d depth reconstruction from a single still image. *IJCV* (2008)
27. Sharon, E., Galun, M., Sharon, D., Basri, R., Brandt, A.: Hierarchy and adaptivity in segmenting visual scenes. *Nature* **442** (2006) 810–813
28. Shi, J., Malik, J.: Normalized cuts and image segmentation. *PAMI* (2000)
29. Vecera, S.P., Vogel, E.K., Woodman, G.F.: Lower region: A new cue for figure-ground assignment. *Journal of Experimental Psychology: General* (2002)
30. Vese, L.A., Chan, T.F.: A multiphase level set framework for image segmentation using the Mumford and Shah model. *IJCV* (2002)
31. Yu, S.X.: Angular embedding: from jarring intensity differences to perceived luminance. *CVPR* (2009)
32. Yu, S.X., Lee, T.S., Kanade, T.: A hierarchical Markov random field model for figure-ground segregation. *CVPR* (2001)
33. Yu, S.X., Shi, J.: Segmentation with pairwise attraction and repulsion. *ICCV* (2001)
34. Zhou, H., Friedman, H.S., von der Heydt, R.: Coding of border ownership in monkey visual cortex. *Journal of Neuroscience* (2000)