

Investigating How and When Perceptual Organization Cues Improve Boundary Detection in Natural Images

Leandro A. Loss, George Bebis, Mircea Nicolescu
Computer Vision Laboratory
University of Nevada, Reno
loss,bebis,mircea@cse.unr.edu

Alexei Skurikhin
Space and Remote Sensing Sciences
Los Alamos National Laboratory
alexei@lanl.gov

Abstract

Boundary detection in natural images represents an important but also challenging problem in computer vision. Motivated by studies in psychophysics claiming that humans use multiple cues for segmentation, several promising methods have been proposed which perform boundary detection by optimally combining local image measurements such as color, texture, and brightness. Very interesting results have been reported by applying these methods on challenging datasets such as the Berkeley segmentation benchmark. Although combining different cues for boundary detection has been shown to outperform methods using a single cue, results can be further improved by integrating perceptual organization cues with the boundary detection process. The main goal of this study is to investigate how and when perceptual organization cues improve boundary detection in natural images. In this context, we investigate the idea of integrating with segmentation the Iterative Multi-Scale Tensor Voting (IMSTV), a variant of Tensor Voting (TV) that performs perceptual grouping by analyzing information at multiple-scales and removing background clutter in an iterative fashion, preserving salient, organized structures. The key idea is to use IMSTV to post-process the boundary posterior probability (PB) map produced by segmentation algorithms. Detailed analysis of our experimental results reveals how and when perceptual organization cues are likely to improve or degrade boundary detection. In particular, we show that using perceptual grouping as a post-processing step improves boundary detection in 84% of the grayscale test images in the Berkeley segmentation dataset.

1. Introduction

High quality image segmentation is known to be an essential part of a broad range of computer vision applications including target detection and object recognition. A considerable number of segmentation methods rely on

good boundary detection. Although many different boundary detection methods have been proposed in the literature over the last thirty years, the lack of standard benchmarks has made it difficult to assess their applicability and efficiency in real problems. Recently, the Berkeley segmentation dataset [7] has become an important benchmark for testing modern boundary detection algorithms. It contains a large number of challenging images along with ground truth segmentations (i.e., human labeled images) for assessing accuracy and robustness. Several methods have been evaluated on this dataset and results along with ranking information for each of them are publicly available¹

Traditionally, boundary detection has been performed using a single cue such as brightness or texture. Motivated by studies in psychophysics, Martin et al. [7] have recently proposed a new algorithm that optimally combines information from multiple cues. The main idea is to combine local measurements based on color, texture and brightness using a classifier that was trained using human segmented images. The output of the classifier is the probability that a given pixel lies on a boundary. By applying the classifier at each image location and orientation, they obtain a PB map which is then thresholded to obtain a set of boundaries. Very promising segmentation results have been reported on the Berkeley benchmark using this approach.

Comparing human performance to any of the methods that have reported segmentation results on the Berkeley dataset, however, still shows a significant performance gap. Given that ground truth was obtained in an unconstrained way, it is likely that humans used many other aspects of human perception to obtain segmentations that extrapolate the scope of modern algorithms. For example, humans might have used high-level image understanding and object recognition skills to produce more meaningful segmentations while most computer applications depend on image segmentation and boundary detection to achieve some image understanding or object recognition. Therefore, it would

¹<http://www.eecs.berkeley.edu/Research/Projects/CS/vision/grouping/segbench/>

be unlikely that single-aspect segmentation algorithms (i.e., employing only photometric information) would be capable of achieving close to human performance.

A weakness of the method proposed by Martin et al. [7] is that no perceptual organization cues were employed in the segmentation process. Although they admit explicitly that perceptual organization cues were ignored intentionally in their approach, they also note that the PB map produced by their method could become an integral part to any perceptual organization algorithm that operates on natural images, either by grouping regions or edge segments. In a later study by the same group [15], mid-level cues (e.g., continuity) were considered for improving boundary detection.

Our goal in this paper is to investigate how and when perceptual organization cues are likely to improve boundary detection in natural images. In this context, we propose integrating the IMSTV, a variant of TV that performs perceptual grouping by analyzing information at multiple-scales and removing background clutter in an iterative fashion, preserving salient, organized structures. The key idea is to use IMSTV to post-process the PB map produced by segmentation algorithms. IMSTV produces a new PB map by integrating perceptual organization cues which is then used to obtain the final boundaries. It should be mentioned that although we have adopted IMSTV in this study, other perceptual grouping algorithms might also be appropriate for post-processing the results of segmentation.

For evaluation purposes, we have experimented with grayscale images considering five different segmentation methods from the Berkeley segmentation benchmark: Brightness Gradient (BG), Gradient Magnitude (GM), Multi-Scale Gradient Magnitude (MGM), Texture Gradient (TG), and Brightness/Texture Gradients (BTG) [7]. These methods were chosen among others because they are among the top performers in the benchmark. Nevertheless, the proposed framework is general enough to accommodate any other segmentation method. Detailed analysis of our experimental results provides useful insight revealing how and when perceptual organization cues are likely to improve or degrade boundary detection results. Overall, IMSTV appears to be a valuable post-processing tool which improves the results of boundary detection.

2. Boundary Detection Methods

In this section, we briefly review the five boundary detection methods integrated with IMSTV: BG, GM, MGM, TG, and BTG. The first four methods perform boundary detection using a single cue while the last method combines information from two different cues using the method of Martin et al. [7]. Each method produces a PB map which is used as input to IMSTV. IMSTV outputs a new PB map by incorporating perceptual organization cues. A PB map can be visualized as an image whose pixel intensity encodes the

probability that a pixel lies on a boundary. The higher the pixel intensity, the higher the probability that the pixel lies on a boundary. Thresholding the PB map yields a set of boundaries in an image.

A common characteristic to all five methods is their reliance on image photometric information to build a PB map. The GM method computes image gradient magnitudes at each pixel to produce the PB map. The gradients are estimated using a pair of Gaussian derivative filters at a unique, learned, optimal scale. Learning was performed using 200 training images from the Berkeley segmentation dataset. The MGM method computes image gradient magnitudes at two different scales to produce the PB map. The gradients are estimated at each pixel using pairs of Gaussian derivative filters at two, also learned, optimal scales. The BG method uses local brightness gradients to obtain the PB map. The gradients are estimated using a χ^2 difference in the distribution of pixel luminance values of two half discs centered at a given pixel and divided in half at the assumed boundary orientation.

The TG method uses local texture gradients to produce the PB map. The gradients are estimated using a χ^2 difference in the distribution of *textons* of two half discs centered at a given pixel and divided in half at the assumed boundary orientation. Textons are computed by clustering the responses of a bank of filters using K-means. The bank of filters was composed of standard even- and odd-symmetric quadrature pair elongated linear filters. The BTG method combines local brightness and texture gradients to obtain the PB. BTG has demonstrated one of the best performances to date on the Berkeley segmentation benchmark. Additional information about each of these methods can be found in [7].

Fig. 1(b-e) show the PB map computed by these methods for the image in Fig. 1(a). The ground truth obtained by five human subjects is shown in Fig. 1(f). All five methods above have been previously evaluated on the Berkeley dataset and represent some of the top performers to date. The PB maps, specific results and ranking information for each method are publicly available from the Berkeley benchmark website.

3. Perceptual Organization and Boundary Detection

Perceptual organization can be defined as the ability to detect organized structures or patterns in the presence of missing and noisy information. Local contour characteristics, such as continuity, symmetry, convexity and parallelism, are representative examples of perceptual organization cues which can be used to reveal important information about the global organization of structures in an image. The importance of perceptual organization cues in human perception has been known for over 90 years by psychologists




Figure 1. The PB map computed by the methods tested in our study: (a) original image, (b) GM PB map, (c) MGM PB map, (d) TG PB map, (e) BGT PB map, (f) ground truth.

[10, 11]. A recent study, summarizing the role of perceptual organization in image segmentation and boundary detection, can be found in [3].

Many boundary detection and segmentation methods have taken advantage of perceptual organization cues, implementing and combining them in different ways. Representative examples of this diversity can be found in [13, 4, 12, 1, 14, 8, 5]. Among these methods, TV represents a general framework which was formalized in [8] and was shown to work well in many practical applications. TV uses a voting process to infer salient structures from noisy and sparse data. It works by representing input data as tensors and interrelating them through voting fields built from a saliency function that incorporates the Gestalt principles of proximity and good continuation. In this study, we have chosen IMSTV [5], a TV variant scheme, in order to investigate the idea of integrating segmentation with perceptual grouping to improve boundary detection. IMSTV was shown to deal better than traditional TV with clutter and objects appearing in different scales in natural images [5, 6]. Instead of our choice, other perceptual grouping schemes might be used for the same purpose.

In IMSTV, like in TV, edge segments are encoded as tensors and communicate to one another through vote casting. The votes are also tensors, whose length and direction are computed through an expression that includes the perceptual cues of proximity and good continuation [8]. The voting process deforms the original tensors revealing salient and organized structures due to the local perceptual agreement of all segments of the same structure with each other, as shown in Fig. 2. Conversely, clutter or noisy segments

are also revealed by not agreeing perceptually within their neighborhood.




Figure 2. Tensor voting process: (a) input points, (b) tensor encoding, (c) deformation of tensors reveals the salient curve.

IMSTV is different from TV in that it performs figure-ground segmentation by analyzing saliency information at multiple scales and removing low-saliency segments in an iterative fashion. In contrast to traditional TV that uses hard thresholding and single-scale analysis, IMSTV removes noisy segments conservatively according to their behavior across a range of scales. Then, it applies re-voting on the remaining segments to estimate their saliency more reliably. These improvements were shown to better handle low signal-to-noise ratio images, as shown in the example provided in Fig. 3. It should be mentioned that IMSTV can be implemented efficiently without having to recompute the votes at each scale and each iteration. The idea is updating the votes in each iteration instead of recomputing them from scratch.




Figure 3. Examples produced by IMSTV on low signal-to-noise ratio images (from [5]).

4. Improving Boundary Detection in Natural Images

In this section, we investigate under what conditions one would expect perceptual organization to improve boundary detection. As mentioned earlier, it is well known that perceptual organization cues play an important role in human perception. Therefore, would it be reasonable to expect that perceptual organization cues would always improve boundary detection results produced by a machine? The answer to this question depends on how we evaluate boundary detection results and, in particular, what kind of ground truth information is used to evaluate performance.

To the best of our knowledge, there is no widely accepted benchmark containing ground truth information specifically obtained for the purpose of evaluating perceptual organization algorithms. Although perceptual organization cues have been employed with relative success in the literature (e.g., see [14, 6]), the results of these studies were never thoroughly compared to human performance specifically addressing perceptual organization cues. Existing benchmarks, such as the Berkeley segmentation benchmark, contain human segmented data. However, it is very unlikely that humans produced these segmentations using perceptual organization cues alone. A more plausible scenario is that they used both perceptual organization and higher-level processes such as object recognition and image understanding. Therefore, these benchmarks might not be the most appropriate in evaluating algorithms employing perceptual organization cues. Nevertheless, in the absence of a more appropriate dataset, we will be using the Berkeley benchmark in our study. However, we should keep in mind the above issues when it comes to evaluating boundary detection results using photometric or perceptual organization cues.

Without intending to exceed or even get close to human performance, we believe that boundary detection methods using photometric cues (i.e., [7]) can be improved by incorporating perceptual organization cues in the detection process. To get a better insight, we have analyzed below certain local configurations in natural images. This analysis can reveal upfront situations where perceptual grouping would be most beneficial, and others where it would be expected to make no improvements or even degrade the results. Let us consider Fig. 4, for example. The regions within the red square in each of the images shown in Fig. 4 (a, b) have been magnified for clarity and shown in Fig. 4 (a.1, b.1). The respective BG and BTG PB maps are shown in Fig. 4 (a.2, b.2), where lighter intensities correspond to a lower probability. One can notice that parts of the contour around the main objects in each image are diminished due to the low contrast between them and the background. However, let us suppose now that we encode these values as tensors, as shown in Fig. 4 (a.3, b.3), where tensor size is given by the PB map value and tensor direction by the normal to the edge direction. If we apply IMSTV, these same contours can be intensified as shown in Fig. 4 (a.4, b.4). This is because the communication between neighboring segments reveals the locally organized structure underlying those contours. In other words, a plausible continuation between the penguin's neck and chest, as well as between the sail's parts, can be found, improving the results produced by BG and BTG.

On the other hand, let us consider Fig. 5. The regions shown by the red squares in each of the images in Fig. 5 (a, b) have been magnified for clarity and shown in Fig. 5 (a.1, b.1). Fig. 5 (a.2, b.2) show the respective GM and BG




Figure 4. Examples illustrating cases where perceptual grouping improves boundary detection (see text for details): (a, b) original images from Berkeley dataset, (a.1, b.1) region within the red squares magnified, (a.2, b.2) BG and BTG PB maps, (a.3, b.3) gradients encoded as tensors in IMSTV, (a.4, b.4) tensors after iterative voting using IMSTV.

PB maps. It should be noted in these cases that GM and BG produced strong responses due to the high contrast between the object and the background. If we encode these values as tensors, as shown in Fig. 5 (a.3, b.3), and apply IMSTV, then these contours will be deteriorated as shown in Fig. 5 (a.4, b.4). This is because the communication between neighboring segments from the jagged edges is weak, since they do not satisfy the rules of good continuation and smoothness. This degrades the results produced by GM and BG.

5. Experimental Results

In order to evaluate the contribution of perceptual organization in boundary detection, we used the IMSTV scheme to post-process the PB map produced by each of the five methods reviewed in Section 2. The output of IMSTV is a new PB map which can be thresholded to obtain a set of boundaries. For evaluation, we used the grayscale test images and the corresponding PB maps from the Berkeley segmentation benchmark. As discussed earlier, pixels in the PB map were encoded as tensors whose size was given by the PB intensity and direction by the normal to the edge direction crossing the pixel (i.e., see Fig. 4(a.3, b.3) and 5(a.3, b.3)).

We initiated the IMSTV voting process at 10 equidistant scales, corresponding to voting fields ranging from 10% to 100% of the smallest image size (i.e., 320 for the Berkeley dataset images), as suggested in [6]. In each iteration, low-saliency segments (i.e., with saliency below 5% of the max-




Figure 5. Examples where perceptual grouping degrades boundary detection (see text for details): (a) original images from the Berkeley dataset, (a.1, b.1) regions within the red squares magnified, (a.2, b.2) GM and BG PB maps, (a.3, b.3) magnitudes encoded as tensors in IMSTV, (a.4, b.4) tensors after iterative voting using IMSTV.

imum saliency) were eliminated and voting was repeated using the remaining tensors. This process continues, removing low-salient segments for a pre-determined number of iterations (i.e., 10 in our experiments). In [6], several different stopping criteria were suggested to avoid over-segmentation. However, we decided not to implement them here in order to generate uniform PB maps. The way we build a PB map for IMSTV is by counting the number of iterations each pixel survived the elimination process. The longer a pixel is retained, the higher is its probability to belong to an organized structure in the image.

To quantify boundary detection results, we used Precision-Recall Curves (PRCs) like in the Berkeley segmentation benchmark. PRCs reflect the trade-off between true boundary pixels detected and non-boundary pixels detected at a given threshold. It should be mentioned, however, that all comparisons in the Berkeley benchmark were carried out using the F-measure [9], which is a weighted harmonic mean of precision (P) and recall (R): $F = PR / (\alpha R + (1 - \alpha)P)$ where α is a weight. The value of α was set to .5 in [7] which is usually called the *equal regime*. Different values of α allow for different regimes (e.g., *high precision regime* for $\alpha > .5$, or *high recall regime* for $\alpha < .5$).

To avoid any bias towards a specific regime and evaluate overall performance more objectively, we have also computed the Area Above the precision-recall Curve (AAC) in our experiments. The use of AAC's dual, the Area Under a Curve (AUC), has been investigated in other studies (e.g.,

[2]), suggesting that AUC is a better measure for evaluating overall performance instead of using a single measurement on the curve. In our case, our objective is minimizing AAC in order to improve both precision and recall rates.

Fig. 6 shows the PRCs for each of the five boundary detection methods tested. Each graph also shows the corresponding PRC using IMSTV for post-processing. Each curve is the average over 100 PRCs corresponding to the 100 grayscale test images in the Berkeley segmentation dataset. Table 1 shows the F-measure and AAC values for each PRC. As it can be noted, at equal regime, perceptual grouping partially improved only one method (i.e., GM), slightly degraded one method (i.e., TG), and partially improved or degraded the rest (i.e., MGM, BG, and BTG). Considering the AAC measure, however, IMSTV improved all methods except TG. The reason that TG was not improved is because most boundaries found using texture gradient violate the perceptual organization rules used by IMSTV. For the methods with shown improvement, it is interesting to note that post-processing improved the results at certain thresholds, that is, more improvements can be noticed at a high precision regime.

Table 1. Resulting F-measure (F) at equal regime and AAC for the five methods tested with and without IMSTV.

| Method | Original | | w/ IMSTV | |
|--------|----------|-----|----------|-----|
| | F | AAC | F | AAC |
| GM | .56 | .43 | .57 | .38 |
| MGM | .58 | .31 | .58 | .28 |
| TG | .58 | .21 | .57 | .24 |
| BG | .60 | .34 | .60 | .31 |
| BTG | .63 | .28 | .62 | .26 |

Looking at the PRCs alone does not provide sufficient information to appreciate the benefits of integrating perceptual organization cues with segmentation. Tables 2 and 3 provide more information to further analyze the results obtained. Specifically, each table shows the actual Number of Images Improved (NII) after post-processing, the Average Improvement Rate (AIR), the Number of Images Degraded (NID) after post-processing, and the Average Degradation Rate (ADR) for each method. Table 2 shows these statistics using the F-measure while Table 3 shows the same statistics using the AAC value. The results based on the F-measure indicate that although the number of images improved is lower than the number of images degraded, the average rate of improvement is usually higher than the average rate of degradation. In other words, the rate of improvement is higher for the images improved than the rate of degradation for the images damaged. Considering the same statistics in the case of AAC, it is more clear that IMSTV is really beneficial as a post-processing step. It has not only improved more images, the rate of improvement is also higher on the




Figure 6. Average PRCs comparing each method with and without IMSTV to post-process boundary detection results: (a) GM, and GM with IMSTV, (b) MGM, and MGM with IMSTV, (c) TG, and TG with IMSTV, (d) BG, and BG with IMSTV, (e) BTG, and BTG with IMSTV. The resulting F-measure and AAC are shown in Table 1.

average. At the same time, it has degraded fewer images with a lower rate on the average.

Table 2. Results based on the F-measure at equal regime obtained by post-processing the 100 test images from the Berkeley dataset using IMSTV.

| Method | NII | AIR | NID | ADR |
|--------|-----|------|-----|------|
| GM | 40 | 9.0% | 60 | 5.3% |
| MGM | 35 | 5.5% | 65 | 3.7% |
| TG | 24 | 3.1% | 76 | 4.1% |
| BG | 40 | 4.9% | 60 | 4.9% |
| BTG | 36 | 4.4% | 64 | 3.4% |

Table 3. Results based on AAC obtained by post-processing the 100 test images from the Berkeley dataset using IMSTV.

| Method | NII | AIR | NID | ADR |
|--------|-----|------|-----|------|
| GM | 71 | 8.4% | 29 | 5.2% |
| MGM | 72 | 5.2% | 28 | 3.8% |
| TG | 62 | 4.1% | 38 | 3.5% |
| BG | 82 | 6.9% | 18 | 3.6% |
| BTG | 84 | 7.2% | 16 | 3.6% |

A detailed analysis of these results can reveal even more

information about the kind of images that are more likely to be improved by IMSTV. Table 4 shows the number of images improved by IMSTV, considering the F-measure at equal regime, relative to the F-measure obtained by the original methods. The results show that 53% to 87.5% of the images resulting in F-measures originally below .5 were improved. As the resulting F-measure increases, the rate of improved images decreases. These results indicate that perceptual organization cues are especially beneficial to images having low F-measures. Although we would have to experiment more to further verify this observation, it appears that such images are not well explained by the features extracted (e.g., Fig. 4). On the other hand, when the features extracted can explain an image well (e.g., Fig. 5), then post-processing seems to have less effect.

Table 5 shows the number of images improved by IMSTV, considering the AAC value relative to the F-measure obtained by the original methods. Although 70.0% to 92.3% of the images resulting in F-measures originally below .5 were improved, it is interesting to note that, in general, high rates were achieved throughout the whole F-measure range. These results suggest that independently of the performance achieved by a given method, it might be

Table 4. Improvement based on the F-measure at equal regime relative to the original F-measure.

| Method | [.0,.5] | (.5,.6] | (.6,.7] | (.7,.8] | (.8,1.] |
|--------|---------|---------|---------|---------|---------|
| GM | 87.5% | 55.6% | 30.0% | 8.7% | 0.0% |
| MGM | 75.0% | 51.9% | 23.5% | 5.0% | 0.0% |
| TG | 53.3% | 26.0% | 15.6% | 20.0% | 0.0% |
| BG | 84.6% | 69.6% | 39.3% | 6.5% | 0.0% |
| BTG | 70.0% | 70.6% | 40.6% | 11.4% | 0.0% |

always possible to improve its overall performance using perceptual organization cues for post-processing.

Fig. 7 shows some boundary detection results for each method with and without IMSTV. As it can be observed, IMSTV eliminates noisy segments more effectively, preserving boundary segments that satisfy the perceptual organization principles underlying IMSTV.

Table 5. Improvement based on the AAC relative to the original F-measure.

| Method | [.0,.5] | (.5,.6] | (.6,.7] | (.7,.8] | (.8,1.] |
|--------|---------|---------|---------|---------|---------|
| GM | 87.5% | 92.6% | 60.0% | 52.2% | 50.0% |
| MGM | 87.5% | 85.2% | 58.8% | 70.0% | 33.3% |
| TG | 86.7% | 59.2% | 57.8% | 60.0% | 33.3% |
| BG | 92.3% | 95.7% | 67.9% | 77.4% | 100.0% |
| BTG | 70.0% | 100.0% | 84.4% | 82.9% | 66.7% |

6. Conclusions and Future Work

We investigated the idea of integrating perceptual organization cues with segmentation in order to improve boundary detection in natural images. In this context, we proposed using IMSTV, a TV variant, to post-process the PB map produced by segmentation methods based on one or more photometric cues. For evaluation and comparison purposes, we used the Berkeley segmentation benchmark along with five segmentation methods that have shown good results on this benchmark. Detailed analysis of our results revealed how and when perceptual organization cues are more likely to improve or degrade boundary detection. In particular, our results indicate that IMSTV improved up to 40% of the test images, when considering the F-measure at equal regime as a performance measure. These improvements were especially noticed among images having low F-measures originally, although, in general, a higher performance is more obvious at high precision regime. When considering AAC, IMSTV improved up to 84% of the images and across the entire range of original F-measure. These results look particularly interesting and encouraging to us. The benefits of integrating perceptual organization cues with segmentation methods are quite clear. For future research, we first plan to consider the problem of choosing between photometric features and perceptual organization cues more wisely. This will potentially prevent some of the degrading cases

from happening and, consequently, improve overall results even more. Second, we plan to consider polarity information (first-order voting) in order to better preserve junctions. This will certainly improve the segmentation quality and accuracy.

References

- [1] T. Alter and R. Basri. Extracting salient contours from images: An analysis of the saliency network. *Proc. IEEE Conference on Computer Vision and Pattern Recognition - CVPR'96*, pages 13–20, 1996.
- [2] A. Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159, 1997.
- [3] C. C. Fowlkes, D. R. Martin, and J. Malik. Local figure-ground cues are valid for natural images. *Journal of Vision*, 7(8):1–9, 2007.
- [4] L. Héroult and R. Horaud. Figure-ground discrimination: A combinatorial optimization approach. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 15:899–914, 1993.
- [5] L. Loss, G. Bebis, M. Nicolescu, and A. Skourikhine. Perceptual grouping based on an iterative multi-scale tensor voting. *Proceedings of the 2nd International Symposium on Visual Computing (ISVC06), LNCS*, 4292:1786–1797, 2006.
- [6] L. Loss, G. Bebis, M. Nicolescu, and A. Skourikhine. An automatic framework for figure-ground segmentation in cluttered backgrounds. *Proceedings of the British Machine Vision Conference 2007 (BMVC07)*, 1:202–211, 2007.
- [7] D. R. Martin, C. C. Fowlkes, and J. Malik. Learning to detect natural image boundaries using local brightness, color, and texture cues. *IEEE Transactions on Pattern Recognition and Machine Intelligence*, 26(5):530–549, 2004.
- [8] G. Medioni, M.-S. Lee, and C.-K. Tang. *A Computational Framework for Segmentation and Grouping*. Elsevier Science, 2000.
- [9] C. V. Rijsbergen. Information retrieval. *second edition, Dept. of Computer Science, University of Glasgow*, 1979.
- [10] E. Rubin. Visuaell wahrgenommene figuren. In *Copenhagen: Gyldendalske*, 1921.
- [11] E. Rubin. Figure and ground. In *In Visual Perception: Essential Readings*, S. Yantis, ed. (Philadelphia: Psychology Press), pages 225–229, 2001.
- [12] S. Sarkar and K. Boyer. Quantitative measures of change based on feature organization: Eigenvalues and eigenvectors. *Proc. IEEE Conference on Computer Vision and Pattern Recognition - CVPR'96*, pages 478–483, 1996.
- [13] S. Ullman and A. Sha'ashua. Structural saliency: The detection of globally salient structures using a locally connect network. *2nd. International Conference on Computer Vision - ICCV'88*, 1988.
- [14] L. Williams and K. Thornber. A comparison measures for detecting natural shapes in cluttered background. *International Journal of Computer Vision*, 34(2/3):81–96, 2000.
- [15] C. F. X. Ren and J. Malik. Mid-level cues improve boundary detection. In *Berkeley Technical Report 05-1382*, 2005.



Figure 7. Visual comparison of results: (a) original grayscale images, (b) initial boundaries detected, (c) resulted boundaries by thresholding at the optimal F-measure (d) resulted boundaries using post-processing, thresholded at the optimal F-measure, (e) ground truth.