Unsupervised Image Partitioning with Semidefinite Programming

Jens Keuchel, Christoph Schnörr, Christian Schellewald, Daniel Cremers

CVGPR-Group, Dept. Math. and Comp. Science
University of Mannheim, D-68131 Mannheim, Germany
Email: {jkeuchel,schnoerr,cschelle,cremers}@ti.uni-mannheim.de
http://www.cvgpr.uni-mannheim.de

Abstract. We apply a novel optimization technique, semidefinite programming, to the unsupervised partitioning of images. Representing images by graphs which encode pairwise (dis)similarities of local image features, a partition of the image into coherent groups is computed by determining optimal balanced graph cuts. Unlike recent work in the literature, we do not make any assumption concerning the objective criterion like metric pairwise interactions, for example. Moreover, no tuning parameter is necessary to compute the solution. We prove that, from the optimization point of view, our approach cannot perform worse than spectral relaxation approaches which, conversely, may completely fail for the unsupervised choice of the eigenvector threshold.

1 Introduction

Partitioning images in an unsupervised way is a common goal of many low-level computer vision applications. Based on some locally computed features like color, texture, or motion, the image should be split into coherent groups whose members look "similar". As no prototypes for the different groups are given in advance, the "correct" partitioning cannot be easily defined. To this end, a hierarchical bi-partitioning approach is often pursued in practice: The image is split into two main parts, which could be split further in subsequent applications of the algorithm. Figure 1 shows two images taken from the VisTex-database [1] that should give you an impression of the difficulty of the partitioning task.

To guide the search for a "good" segmentation, an appropriate representation of the image is as well needed as optimization criteria which give measures of the quality of a segmentation. To this end, the representation of images by graph structures has recently attracted the interest of researchers [2–4]: An image is represented by a graph with locally extracted image features as vertices and pairwise (dis)similarity values as edge weights. The goal is to find a cut through this graph that divides it in two coherent parts. Several methods from spectral graph theory were proposed in the literature to solve this problem [5, 6, 2]. A major problem of these approaches concerns the appropriate choice of a threshold value to split the computed eigenvector in two reasonable parts.

In this paper, we investigate the application of a novel optimization technique, *semidefinite programming (SDP)*, to the field of unsupervised partition-

Fig. 1. A color scene (left) and a gray-value scene comprising some natural textures (right). How to partition such scenes into coherent groups in an *unsupervised* way based on pairwise (dis)similarities between local measurements?

ing. Therefore, by using the graph representation, we derive a problem formulation that yields a quadratic functional defined over binary decision variables which has to be minimized subject to linear constraints. The combinatorial complexity of this optimization task is then dealt with in two steps: Firstly, the decision variables are lifted to a higher-dimensional space where the optimization problem is relaxed to a *convex* optimization problem [7]. Secondly, the decision variables are recovered from the global optimum of the relaxed problem by using a small set of randomly computed hyperplanes [8].

In contrast to related work [3, 4], *no specific assumptions* are made with respect to the functional form apart from a symmetry condition. As a consequence, our approach can also be applied to other computer vision tasks like perceptual grouping or image restoration [9]. Other favourable properties of the semidefinite programming approach are:

- As the relaxed problem is convex, the global optimum can be computed.
- Interior-point algorithms [10] allow to find the optimum in polynomial time.
- No additional tuning parameters are necessary. This is a significant advantage over alternative optimization approaches [11].
- In contrast to spectral relaxation, no choice of a threshold value is necessary.

In the following, we will apply the semidefinite relaxation approach to binary partitioning problems, and compare the results to spectral relaxation methods.

2 Problem Statement: Binary Combinatorial Optimization for Unsupervised Partitioning

Consider a graph $G(V, E)$ with locally extracted image features as vertices V and pairwise similarity values as edge-weights $w : E \subseteq V \times V \to \mathbb{R}_0^+$. We wish to compute a partition of the set V into two coherent groups $V = S \cup \overline{S}$. Representing such a partition by an indicator vector $x \in \{-1, +1\}^n$, a measure for the partition can be defined as the weight of the corresponding cut:

$$w(S, \overline{S}) = \sum_{i \in S, j \in \overline{S}} w(i, j) = \frac{1}{8} \sum_{i,j \in V} w(i, j)(x_i - x_j)^2 = \frac{1}{4} x^\top L x \; . \qquad (1)$$

Here, $L = D - W$ denotes the Laplacian matrix of the graph G, and D is the diagonal matrix with the entries $d(i,i) = \sum_{j \in V} w(i,j)$. As the weight function w encodes a similarity measure between pairs of features, coherent groups correspond to low values of $w(S, \overline{S})$.

In order to avoid unbalanced partitions which are likely when just minimizing $w(S, \overline{S})$, a classical partitioning approach is to demand that both groups contain the same number of vertices by adding the constraint $e^\top x = 0, e = (1, \ldots, 1)^\top$, hence arriving at the following combinatorial minimization problem:

$$\inf_x x^\top L x \; , \quad e^\top x = 0 \; , \quad x \in \{-1, +1\}^n \; . \tag{2}$$

Since e is the eigenvector of the Laplacian matrix L with eigenvalue 0, a natural relaxation of this problem is to drop the integer constraint and compute the eigenvector corresponding to the second smallest eigenvalue of L (the so-called "Fiedler vector"; see, e.g. [5]). An approximate solution for (2) is then derived by thresholding this eigenvector. This raises the question for an appropriate choice of the threshold value. Two natural approaches seem to be promising: To threshold at 0 (because of the $+1/-1$-constraint on x) or at the median of the eigenvector (to meet the balancing constraint $e^\top x = 0$). However, we will show below that an unsupervised choice of the threshold value may fail completely.

In this paper, we focus on the semidefinite relaxation of (2), which *directly* takes into account the integer constraint with respect to x_i, $i = 1, \ldots, n$, instead of just doing so by thresholding afterwards. We will show that this approach compares favorably to the computation of the Fiedler vector both in theory (Section 4) and in practice (Section 5).

Recently, Shi and Malik [2] suggested another successful approach, which is similar to (2): They use a normalized objective function which finally results in the following problem (*normalized cut*):

$$\inf_x \frac{x^\top L x}{x^\top D x} \; , \quad e^\top D x = 0 \; , \quad x \in \{-b, 1\}^n \; , \tag{3}$$

where the number b is not known beforehand. Dropping the integer constraint yields a relaxation of (3) which then can be solved by calculating the second smallest eigenvalue of the normalized Laplacian matrix $\tilde{L} = D^{-1/2} L D^{-1/2}$. A relation of this approach to the semidefinite relaxation approach can be derived by replacing the vector e in (2) by the vector De. As the details of this relation are not straightforward (due to the normalization of the objective function), they are beyond the scope of this paper and will be reported elsewhere [12].

Note that the positivity of the edge-weights w is essential for both spectral relaxation methods; however, the semidefinite relaxation described in the next section only requires the matrix L in (2) to be symmetric, and thus can also be applied in the case of *non-positive* edge-weigths!

3 Semidefinite Relaxation

To derive the semidefinite relaxation of (2), we first replace the linear and the integer constraint, respectively, by quadratic ones: $(e^\top x)^2 = 0 \; , x_i^2 - 1 = 0, i =$

$1, \ldots, n$. Denoting the Lagrangian multiplier variables with $y_i, i = 0, \ldots, n$, the corresponding Lagrangian of (2) yields the following relaxed problem formulation of (2) after some standard transformations:

$$z_d := \sup_{y_0, y} e^\top y \ , \quad L - y_0 ee^\top - D(y) \in \mathcal{S}_+^n \ , \tag{4}$$

where $D(y)$ denotes the diagonal matrix with the vector y on its diagonal, and $\mathcal{S}_+^n$ is the set of positive semidefinite matrices. As this set is a cone (i.e. a special convex set), we arrive at a *convex* optimization problem. The relation to our original problem can be seen by deriving the dual problem of (4):

$$z_p := \inf_{X \in \mathcal{S}_+^n} L \bullet X \ , \quad ee^\top \bullet X = 0 \ , \ D(X) = I \ . \tag{5}$$

Here $L \bullet X = \mathrm{Tr}\,[L^\top X]$ denotes the standard matrix inner product, and $D(X)$ is the matrix X with the off-diagonal elements set to zero. Notice that problem (5) again is convex!

If we rewrite the objective function of (2) as $\inf_x x^\top L x = \inf_x L \bullet xx^\top$, we immediately see the connection to the relaxed problem (5): The rank one matrix $xx^\top$ is replaced by an arbitrary matrix $X \in \mathcal{S}_+^n$, whereas the constraints are just lifted in the higher-dimensional space accordingly.

The elegant duality theory corresponding to the class of convex optimization problems [10] guarantees under mild conditions that optimal primal and dual solutions $X^*, (y_0^*, y^*)$ for (5) and (4) exist and that they yield no duality gap: $z_p - z_d = L \bullet X^* - e^t y^* = 0$.

For the problems considered in this paper, the constraint $ee^\top \bullet X = 0$ requires that the smallest eigenvalue of X is equal to 0, so that no strictly interior point for the primal problem (5) exists. Due to this observation we decided to use the dual-scaling algorithm from [13] for our experiments. This algorithm has the advantage that it does not need to calculate an interior solution for the primal problem during the iterations, but only for the dual problem.

To find a combinatorial solution x_S based on the solution matrix X^* to the convex optimization problem (5), we used the randomized-hyperplane technique proposed in [8]. As this technique does not take into account the constraint $e^\top x = 0$, the solution x_S does not need to be feasible for (2), and the resulting objective value $z_S = x_S^\top L x_S$ may be even smaller than the optimal value of the semidefinite relaxion z_d. Therefore, some modifications of the randomized-hyperplane technique have been proposed in the literature [14, 15]. However, we decided to stick to it as for the applications considered in this paper, it is not mandatory to find a feasible solution: The constraint $e^\top x = 0$ only serves as a strong bias to guide the search to a solution that is balanced reasonably.

4 Comparison to Spectral Relaxation

Poljak and Rendl [16] proved the following relation of the semidefinite relaxation (5) to an eigenvalue optimization problem:

$$z_d = \sup_{v \in \mathbb{R}^n} n \lambda_{\min}\big(V^\top \big(L + D(v)\big) V\big) \ , \quad e^\top v = 0 \ , \tag{6}$$

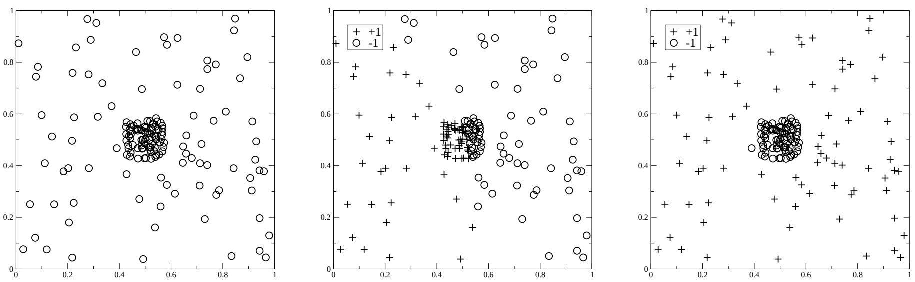

Fig. 2. Point set clustering. The corresponding graph contains a vertex for each point, and the edge weights are calculated from the Euclidian distances $d(i,j)$ between all points by $w(i,j) = \exp(-(\frac{d(i,j)}{0.05})^2)$. **Left:** Input data with 160 points. **Middle:** Solution computed with the Fiedler vector, thresholded at the median value: Spectral relaxation fails! **Right:** Solution computed with SDP.

where $V \in \mathbb{R}^{n\times(n-1)}$ contains an orthonormal basis of the complement $e^\perp$, i.e. $V^\top e = 0$, $V^\top V = I$. An immediate consequence is the following lemma, which shows the connection between the semidefinite relaxation and the computation of the Fiedler vector:

Lemma 1 *The Fiedler vector yields a lower bound for (2) of* $n\lambda_{\min}(V^\top LV)$, *which is a* **weaker** *relaxation than (6):* $n\lambda_{\min}(V^\top LV) \leq z_d$.

It is easy to construct examples where spectral relaxation is too weak and thus cannot compute a meaningful solution. In Figure 2, for example, the Fiedler vector is not able to seperate the dense cluster from the background, despite the fact that the balancing constraint is strictly enforced by the median threshold! In contrast to that, our approach finds two groups of nearly the same size (78 and 82 points).

5 Experiments

The results of the semidefinite relaxation approach for various binary partitioning problems are shown in Figures 3–5, and compared with segmentations obtained with the Fiedler vector. For all experiments, the edge weights $w(i,j)$ building the similarity matrix were computed from the distances $d(i,j)$ between the extracted image features i and j as $w(i,j) = \exp(-\frac{(d(i,j))^2}{\sigma})$, where $d(i,j)$ and σ were chosen application dependent. We studied two different approaches:

(a) Compute $w(i,j)$ for all feature pairs (i,j) directly.

(b) Compute $w(i,j)$ only for neighboring features, and derive the other edge weights by calculation of a path connecting them. This was done by changing the similarity weights to dissimilarities, computing shortest paths, and transforming the weights back afterwards.

For a survey of numerous (dis)similarity measures, see [17].

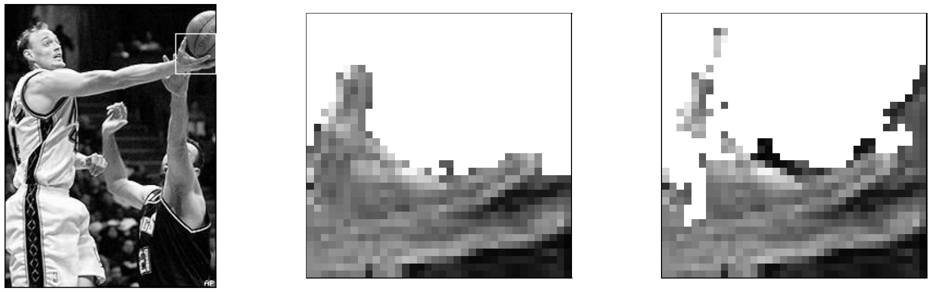

Fig. 3. Grayscale image partitioning. Each pixel is taken as a graph vertex, and the edge weigths are computed from the gray value differences $d(i, j)$ of adjacent pixels with method (b). **Left:** Input image (36×36 pixels) as part of a larger image. **Middle:** Segmentation computed with SDP: The hand is clearly separated from the ball. **Right:** Segmentation computed with the Fiedler vector: No clear separation is obtained by median thresholding. Thresholding at 0 just separates one pixel from the rest of the image.

The results approve the theoretical superiority of the semidefinite relaxation approach: Whereas even a supervised choice of the threshold value for the Fiedler vector does not necessarily yield satisfactory partitionings, the segmentations obtained with SDP are always very reasonable.

6 Conclusion

The results presented in this paper show that the semidefinite relaxation approach is well suited to perform unsupervised binary partitioning for a wide range of applications. It compares favorably both in theory and in practice to the computation of the Fiedler vector. In our further work we will study which other constraints are useful for unsupervised partitioning and could be incorporated into the semidefinite relaxation approach.

Acknowledgement. This work has been supported by the Deutsche Forschungsgemeinschaft (DFG; grant Schn457/3).

References

1. Vision texture database. http://www-white.media.mit.edu/vismod/imagery/VisionTexture/vistex.html.
2. J. Shi and J. Malik. Normalized cuts and image segmentation. *IEEE Trans. Patt. Anal. Mach. Intell.*, 22(8):888–905, 2000.
3. H. Ishikawa. *Global Optimization Using Embedded Graphs.* PhD thesis, Dept. Comp. Science, Courant Inst. Math. Sciences, New York University, 2000.
4. Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Trans. Patt. Anal. Mach. Intell.*, 23(11):1222–1239, 2001.
5. B. Mohar and S. Poljak. Eigenvalues in combinatorial optimization. In R.A. Brualdi et al., editors, *Combinatorial and Graph–Theoretical Problems in Linear Algebra*, volume 50 of *IMA Vol. Math. Appl.*, pages 107–151. Springer, 1993.

Fig. 4. Color image partitioning. We first compute an oversegmentation by applying the mean shift technique [18] at a fine spatial scale in order not to destroy any perceptually significant structure. Instead of thousands of pixels, the graph vertices are then formed by the obtained clusters, and $d(i,j)$ is computed as the color difference of two clusters in the perceptually uniform LUV space. **Top left:** Input image (298×141 pixels), yielding 209 clusters. **Top right:** Segmentation computed with the Fiedler vector thresholded at the median, using method (b). The requirement that both parts have the same size influences the result negatively. **Middle:** Segmentation computed with SDP, using method (b): Spatially coherent structures are favored. **Bottom:** Segmentation computed with SDP, using method (a): Similar colors are grouped together. Notice that for both methods, no choice of a threshold value is necessary for SDP!

6. Y. Weiss. Segmentation using eigenvectors: a unifying view. In *Proc. Int. Conf. Comp. Vision (ICCV'99)*, pages 975–982, 1999.
7. L. Lovász and A. Schrijver. Cones of matrices and set–functions and 0–1 optimization. *SIAM J. Optimization*, 1(2):166–190, 1991.
8. M.X. Goemans and D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42:1115–1145, 1995.
9. C. Schellewald, J. Keuchel, and C. Schnörr. Image labeling and grouping by minimizing linear functionals over cones. In M. Figueiredo et al., editors, *Energy Minimization Methods in Computer Vision and Pattern Recognition*, volume 2134 of *Lect. Not. Comp. Sci.*, pages 235–250. Springer, 2001.
10. Y. Nesterov and A. Nemirovskii. *Interior Point Polynomial Methods in Convex Programming*. SIAM, 1994.

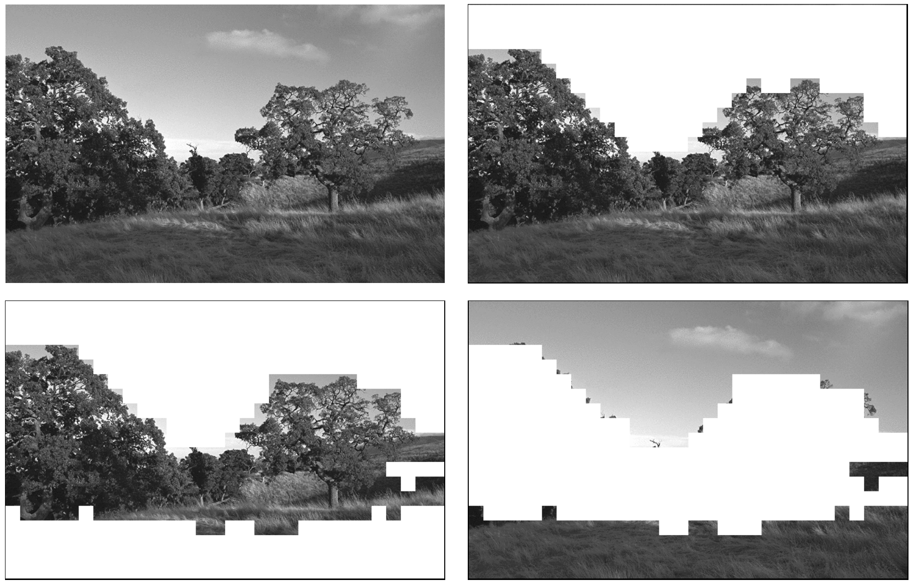

Fig. 5. Grayscale-texture partitioning. The texture measure is derived by subdividing the image into 24×24-pixel windows, and calculating local histograms for two texture features within these windows. Each window corresponds to a graph vertex, and $d(i, j)$ is computed as the χ^2-distance of the histograms using method (a). **Top left:** Input image (720×456 pixels), yielding 570 vertices. **Top right:** Segmentation computed with the Fiedler vector, using the threshold value 0. The median threshold does not make sense here, as the image does not contain two parts of the same size. **Bottom:** Segmentation computed with SDP: Considering the simplicity of this texture measure, the segmentation result is excellent.

11. T. Hofmann and J. Buhmann. Pairwise data clustering by deterministic annealing. *IEEE Trans. Patt. Anal. Mach. Intell.*, 19(1):1–14, 1997.
12. J. Keuchel, C. Schnörr, C. Schellewald, and D. Cremers. Binary partitioning, perceptual grouping, and restoration with semidefinite programming. Technical Report 5/2002, Computer Science Series, CVGPR Group, Univ. Mannheim, 2002.
13. S.J. Benson, Y. Ye, and X. Zhang. Mixed linear and semidefinite programming for combinatorial and quadratic optimization. *Optimiz. Methods and Software*, 11&12:515–544, 1999.
14. A. Frieze and M. Jerrum. Improved approximation algorithms for max k-cut and max bisection. *Algorithmica*, 18:67–81, 1997.
15. Y. Ye. A .699-approximation algorithm for max-bisection. *Mathematical Programming*, 90:101–111, 2001.
16. S. Poljak and F. Rendl. Nonpolyhedral relaxations of graph-bisection problems. *SIAM Journal on Optimization*, 5:467–487, 1995.
17. J. Puzicha, J.M. Buhmann, Y. Rubner, and C. Tomasi. Empirical evaluation of dissimilarity measures for color and texture. In *7th Int. Conf. on Comp. Vision (ICCV'99)*, pages 1165–1172, 1999.
18. D. Comaniciu and P. Meer. Mean shift analysis and applications. In *7th Int. Conf. on Comp. Vision (ICCV'99)*, pages 1197–1203, 1999.