Convex Relaxations for Binary Image Partitioning and Perceptual Grouping

Jens Keuchel, Christian Schellewald, Daniel Cremers, Christoph Schnörr

Computer Vision, Graphics, and Pattern Recognition Group
Department of Mathematics and Computer Science;
University of Mannheim, D-68131 Mannheim, Germany

Email: {jkeuchel,cschelle,cremers,schnoerr}@ti.uni-mannheim.de;
fax: +49 621 181 2744; http://www.cvgpr.uni-mannheim.de

Abstract. We consider approaches to computer vision problems which require the minimization of a global energy functional over binary variables and take into account both local similarity and spatial context. The combinatorial nature of such problems has lead to the design of various approximation algorithms in the past which often involve tuning parameters and tend to get trapped in local minima.
In this context, we present a novel approach to the field of computer vision that amounts to solving a convex relaxation of the original problem without introducing any additional parameters. Numerical ground truth experiments reveal a relative error of the convex minimizer with respect to the global optimum of below 2% on the average.
We apply our approach by discussing two specific problem instances related to image partitioning and perceptual grouping. Numerical experiments illustrate the quality of the approach which, in the partitioning case, compares favorably with established approaches like the ICM-algorithm.

1 Introduction

The minimization of energy functionals plays a central role in many computer vision problems. When additionally discrete decision variables are involved, this task becomes intrinsically combinatorial and hence not easy to tackle. This fact motivated the development of various optimization approaches, like simulated annealing for binary image restoration [1], the ICM-algorithm for Markov Random Field based estimates [2], deterministic annealing for perceptual grouping [3], and many more (see [4]).

However, two crucial requirements from the optimization point-of-view continue to challenge these results: The quality of suboptimal solutions that can be achieved and the presence of additional parameters that have to be tuned by the user. Concerning quality, only simulated annealing can guarantee to find the optimal solution (see [5]) at the cost of unpractically slow annealing schedules. Other approaches are not immune against being trapped in an – eventually bad – local minimum. On the other hand, additional algorithmic tuning parameters for controlling search heuristics, annealing schedules, etc., often critically influence the quality of a solution despite having nothing to do in the first place with the original problem to be solved.

One possibility to overcome this latter point is to consider the mathematically well-understood class of convex optimization problems: They only have one global minimum which can be determined by established numerical algorithms without using additional parameters. In order to tackle a highly non-convex combinatorial optimization problem in this way, it is relaxed by mapping it into a higher dimensional space where the feasible solutions are contained in a convex set (see [6]). Minimization in that space and mapping back yields a suboptimal solution of the original problem, which usually is of a good quality. For special combinatorial problems it is even possible to give an optimality bound: Goemans and Williamson [7] proved the remarkable result that suboptimal solutions computed by convex optimization for the well known max-cut problem are at most 14% worse than the global optimum.

These favorable properties of convex optimization approaches have motivated our work. In the following sections we show how problems from the general class of binary quadratic functionals may be relaxed and solved by convex optimization. We next apply our approach to two specific computer vision problems (binary image partitioning and perceptual grouping) that fit into this class and present numerical results. Additionally, we show the quality of our approach by checking it using one-dimensional signals for which the optimal solution (ground truth) can be easily computed, and by comparing the image partitioning results with those of the ICM-algorithm.

2 Problem Statement: Minimizing Binary Quadratic Functionals

In this paper, we consider the problem of minimizing binary quadratic functionals which have the following general form:

$$J(x) = x^t Q x + 2b^t x + const \ , \quad x \in \{-1,1\}^n,\ Q \in \mathcal{S}^n,\ b \in \mathbb{R}^n \ . \tag{1}$$

As no constraints are imposed on the matrix Q (apart from being in the class of symmetric matrices $\mathcal{S}^n$), such a functional is generally not convex. Furthermore, the integer constraint $x_i \in \{-1,1\}$, $i = 1,\dots,n$, makes the minimization problem (1) intrinsically difficult.

Computer vision problems with quadratic functionals like (1) arise in various contexts; in this paper, we will consider two representatives that are briefly illustrated in the following sections.

2.1 Binary Image Partitioning

Suppose that for each pixel position i of an image, a feature value $g_i \in \mathbb{R}$ has been measured that originates from either of two known prototypical values u_1, u_2[1].

[1] Besides the case of gray-values considered in this paper, local features related to motion, texture, etc., could also be dealt with.

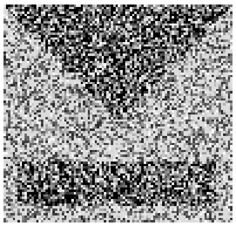

Fig. 1. A binary image, heavily corrupted by (real-valued) noise.

Figure 1 shows an example, where g gives the gray-values of the noisy image that originally was black and white.

To restore the discrete-valued image function x that represents the original image, we wish to minimize the following functional which has the form of (1):

$$J(x) = \frac{1}{4} \sum_i \big((u_2 - u_1) x_i + u_1 + u_2 - 2g_i \big)^2 + \frac{\lambda}{2} \sum_{\langle i,j \rangle} (x_i - x_j)^2 , \; x_i \in \{-1, 1\}, \; \forall i \, . \tag{2}$$

Here, the second term sums over all pairwise adjacent variables in vertical and horizontal direction on the regular image grid.

Functional (2) comprises two terms familiar from many regularization approaches [8]: A data-fitting term and a smoothness term modeling spatial context. However, due to the integer constraint $x_i \in \{-1, 1\}$, the optimization problem considered here is much more difficult than standard regularization problems.

In comparison to the ICM-algorithm [2], which minimizes a similar objective function, there are two main differences: Whereas the functional (2) sums over all pixels and hence is a global approach, the ICM-algorithm optimizes the objective function locally for each pixel in turn, which results in a much smaller complexity. The second difference is that the ICM-algorithm also uses the diagonally adjacent variables in the smoothness term for each pixel.

2.2 Perceptual Grouping

The task in perceptual grouping is to separate familiar configurations from the (unknown) background. Figure 2 shows an example: The instances of a given object consisting of three edge-elements are to be found in a noisy image.

To describe this problem mathematically, let $g_i, i = 1, \ldots, n$, denote a feature primitive (e.g. an edgel computed at location i) in the image plane. We suppose that for each pair of primitives g_i, g_j a similarity measure d_{ij} can be computed corresponding to some given specific object properties (e.g. the difference between the relative angle of the edgels and the expected one according to the given object). Using the spatial context modeled by the values d_{ij}, each primitive g_i is labeled with a decision variable $x_i \in \{-1, 1\}$ ("1" corresponding to

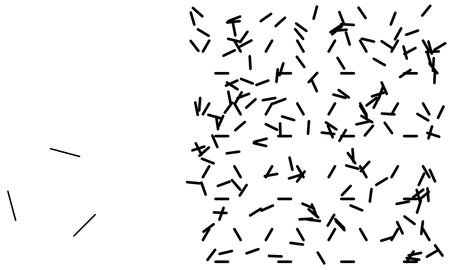

Fig. 2. Left: An object. **Right:** The object was rotated by a fixed arbitrary angle, and translated and scaled copies have been superimposed by noise.

figure, "-1" corresponding to background and noise) by minimizing a functional of the form (1):

$$J(x) = \sum_{\langle i,j \rangle} (\lambda - d_{ij}) x_i x_j + 2 \sum_i (\lambda n - \sum_j d_{ij}) x_i, \ x_i \in \{-1, 1\}, \ \forall i \ , \qquad (3)$$

where the first term sums over all pairs of feature primitives (see [3] for further details).

3 Convex Problem Relaxation and Algorithm

To overcome the hard integer constraints, problem (1) is relaxed to arrive at a less restricted, convex optimization problem of real variables, which can be tackled by using known methods like interior point algorithms that lead to an approximate solution for the original problem.

By dropping the constant and homogenizing the objective function, we rewrite the original functional (1) as follows:

$$x^t Q x + 2 b^t x = \begin{pmatrix} x \\ 1 \end{pmatrix}^t L \begin{pmatrix} x \\ 1 \end{pmatrix} \ , \quad L = \begin{pmatrix} Q & b \\ b^t & 0 \end{pmatrix} \ . \qquad (4)$$

With slight abuse of notation, we denote the vector $\begin{pmatrix} x & 1 \end{pmatrix}^t$ again by x. Then the minimization problem (4) can be rewritten:

$$\inf_{x \in \{-1,1\}^{n+1}} x^t L x = \inf_{x \in \{-1,1\}^{n+1}} L \bullet x x^t \ ,$$

where $X \bullet Y = \mathrm{trace}(XY)$ denotes the standard matrix inner product for two symmetric matrices X, Y.

The convex relaxation is now achieved be replacing the rank-one matrix $xx^t \in \mathcal{K}$ by an arbitrary matrix $X \in \mathcal{K}$, where $\mathcal{K}$ denotes the cone of symmetric, positive semidefinite $(n+1) \times (n+1)$-matrices. Additionally, the integer constraint

$x_i \in \{-1, 1\}$ is *weakly* replaced by the according constraint $x_{ii} = 1$, so that we get:

$$\inf_{X \in \mathcal{K}} L \bullet X \ , \quad x_{ii} = 1, \ \forall i \ . \tag{5}$$

The optimization problem (5) belongs to the class of conic programs, which can be solved approximately with iterative interior-point algorithms (see [9]). To finally get back the solution x to (1) from the computed solution X to (5), we used the randomized-hyperplane technique described in [7].

4 Numerical Experiments

4.1 Restoration of 1D Noisy Signals comprising multiple scales

To investigate the performance of the convex relaxation approach (5), we chose a one-dimensional synthetic signal x' (see Figure 3) that was superimposed with Gaussian white noise with standard deviation $\sigma = 1.0$, and tried to restore the original signal by using the one-dimensional version of functional (2). Since it is possible to compute the global optimum x^* of (1) for one-dimensional signals by using dynamic programming, we were able to compare the resulting suboptimal solution x to the optimal solution x^*. This experiment was repeated 1000 times for different noisy signals and varying values of λ. The results revealed an average relative error of below 2% for both the objective function value and the number of correctly classified pixels (see Figure 4). A representative example of the restoration is given in Figure 5.

4.2 2D-images and Grouping

The results of applying the algorithm to restoration problems in the context of two-dimensional images are shown in Figures 6 – 8. The quality of the restorations using convex optimization is encouraging. Small errors that occur at the corners of connected regions as in Figure 8 (b) are due to the fact that the homogeneity constraint is less restrictive there as the local structure resembles noise.

We also implemented the locally classifying ICM-algorithm (see [2]), and applied it for different values of the parameter $\kappa\beta$, which plays a similar role as the parameter λ in (2). The comparison reveals the power of our global approach: Whereas the ICM-algorithm leaves some local holes within the structure (see Figures 6(d), 7(c), 8(c)), convex optimization is able to close these holes and yield a better piecewise-smooth result.

Concerning the grouping problem, the results of convex optimization are equally good (see Figure 9): Besides a small number of misclassified edgels that couldn't be identified as background by the chosen similarity measure, the original structure is clearly visible.

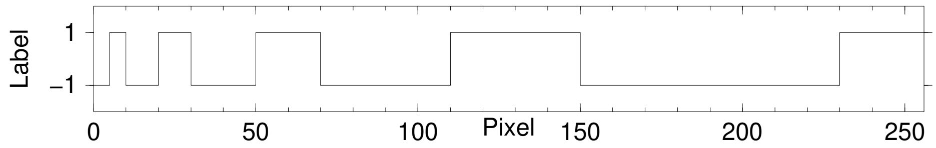

Fig. 3. Signal x' comprising multiple spatial scales.

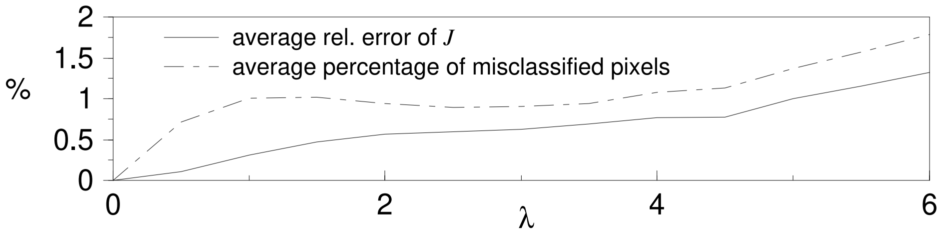

Fig. 4. Average relative error of the objective function J and average percentage of misclassified pixels of the suboptimal solution x compared to the optimal solution x^* for different values of the scale parameter λ.

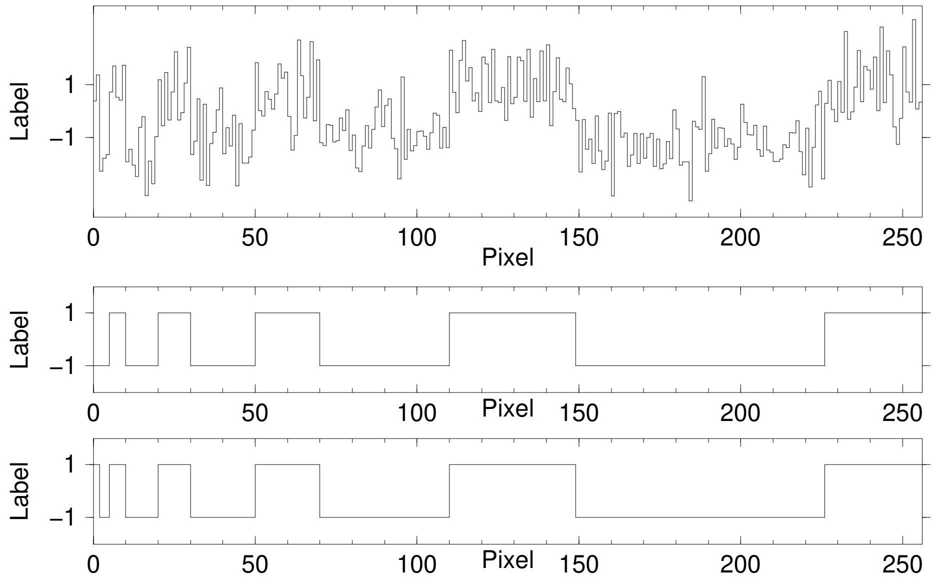

Fig. 5. A representative example illustrating the statistics shown in Fig. 4. **Top:** Noisy input signal. **Middle:** Optimal Solution x^*. **Bottom:** Suboptimal solution x.

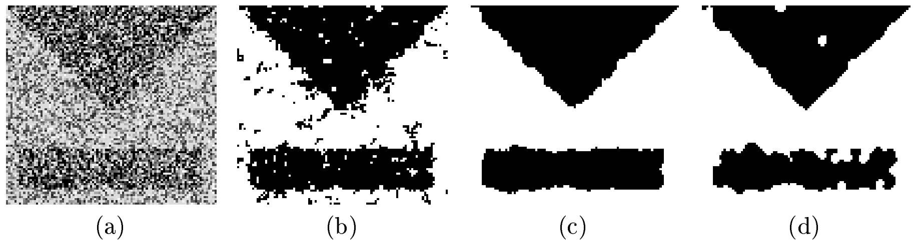

Fig. 6. Arrow and bar real image. **(a)** Noisy original. **(b), (c):** Suboptimal solutions computed with convex optimization for $\lambda = 0.8, 3.0$. **(d)** Solution computed with the ICM-algorithm for $\kappa\beta = 0.99$.

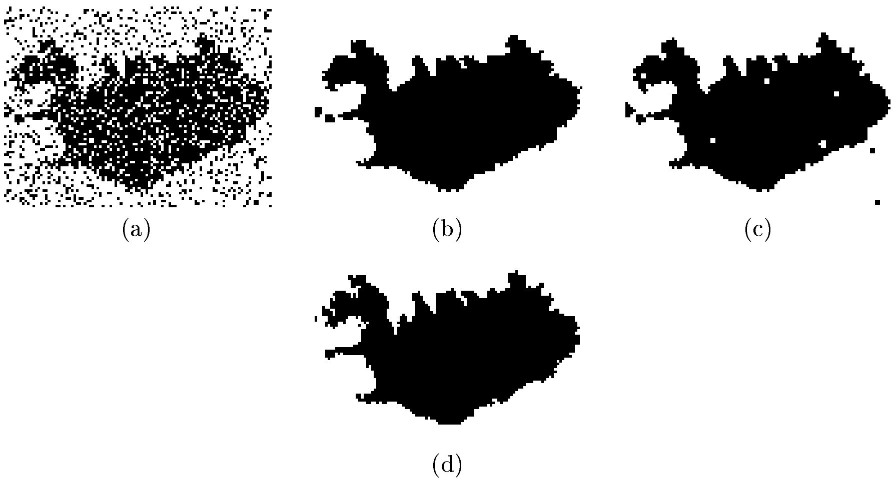

Fig. 7. Iceland image. **(a)** Binary noisy original. **(b)** Suboptimal solution computed with convex optimization for $\lambda = 2.0$. **(c)** Solution computed with the ICM-algorithm for $\kappa\beta = 0.9$. **(d)** Original before adding noise.

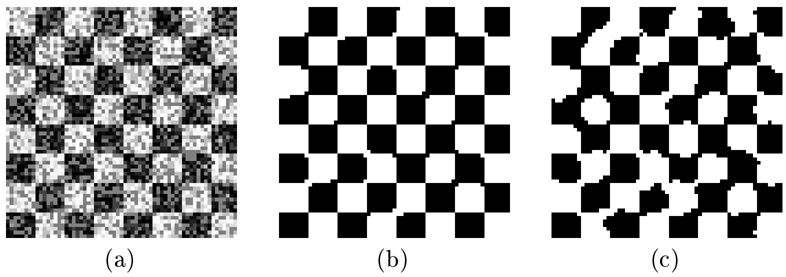

Fig. 8. Checkerboard image. **(a)** Noisy original. **(b)** Suboptimal solution computed with convex optimization for $\lambda = 1.5$. **(c)** Solution computed with the ICM-algorithm for $\kappa\beta = 0.4$.

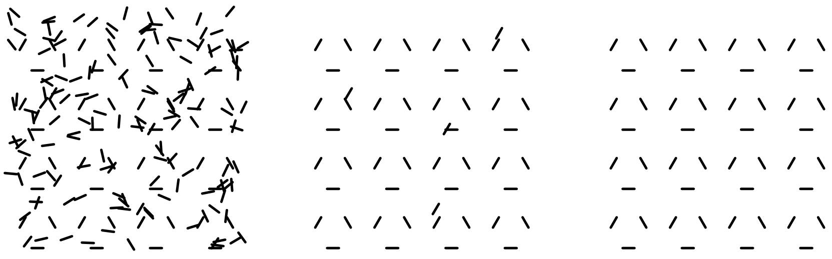

(a) (b) (c)

Fig. 9. Grouping. **(a)** Input data (see Section 2.2). **(b)** Suboptimal solution computed with convex optimization for $\lambda = 0.9$. **(c)** The true solution.

5 Conclusion

The quality of the results of the numerical experiments approves the wide range of applicability of the convex optimization approach. Many other problems with objective functions from the large class of binary quadratic functionals could be tackled. This fact, together with the nice mathematical property that no additional tuning parameters are necessary encourages us to extend our work on the convex optimization approach.

References

1. S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Trans. Patt. Anal. Mach. Intell.*, 6(6):721–741, 1984.
2. J.E. Besag. On the analysis of dirty pictures (with discussion). *J. R. Statist. Soc. B*, 48:259–302, 1986.
3. L. Herault and R. Horaud. Figure-ground discrimination: A combinatorial optimization approach. *IEEE Trans. Patt. Anal. Mach. Intell.*, 15(9):899–914, 1993.
4. C. Schellewald, J. Keuchel, and C. Schnörr. Image labeling and grouping by minimizing linear functionals over cones. In *Proc. Third Int. Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR'01)*, Lect. Notes Comp. Science, INRIA, Sophia Antipolis, France, Sept. 3–5 2001. Springer. *in press.*
5. E. Aarts and J.K. Lenstra, editors. *Local Search in Combinatorial Optimization*, Chichester, 1997. Wiley & Sons.
6. L. Lovász and A. Schrijver. Cones of matrices and set–functions and 0–1 optimization. *SIAM J. Optimization*, 1(2):166–190, 1991.
7. 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.
8. M. Bertero, T. Poggio, and V. Torre. Ill-posed problems in early vision. *Proc. IEEE*, 76:869–889, 1988.
9. 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.