

Performance Evaluation of a Convex Relaxation Approach to the Quadratic Assignment of Relational Object Views

Christian Schellewald, Stefan Roth, Christoph Schnörr

CVGPR–Group
Dept. of Mathematics and Computer Science
University of Mannheim
D-68131 Mannheim, Germany
`{cschelle,roths,schnoerr}@ti.uni-mannheim.de`

Abstract

We introduce a recently published convex relaxation approach for the quadratic assignment problem to the field of computer vision. Due to convexity, a favourable property of this approach is the absence of any tuning parameters and the computation of high-quality combinatorial solutions by solving a mathematically simple optimization problem. Furthermore, the relaxation step always computes a tight lower bound of the objective function and thus can additionally be used as an efficient subroutine of an exact search algorithm.

We report the results of both established benchmark experiments from combinatorial mathematics and random ground-truth experiments using computer-generated graphs. For comparison, a recently published deterministic annealing approach is investigated as well. Both approaches show similarly good performance. In contrast to the convex approach, however, the annealing approach yields no problem relaxation, and four parameters have to be tuned by hand for the annealing algorithm to become competitive.

Keywords: Quadratic assignment, weighted graph matching, combinatorial optimization, convex programming, object recognition

1 Introduction

1.1 Motivation

Visual object recognition is a central problem of computer vision research. A key question in this context is how to represent objects for the purpose of recognition by a computer vision system. Approaches range from view-based to 3D model-based, from object-centered to viewer-centered representations [1], each of which may have advantages under constraints related to specific applications. Psychophysical findings however provide evidence for *view-based* object representations [2] in human vision, and so we will focus on this representation in this paper.




Figure 1: Image features and binary relations computed for an object view using the FEX-system (cf. [3] and <http://www.ipb.uni-bonn.de/ipb/projects/fex/fex.html>). The underlying graph has $|V| = 38$ nodes.

A common and powerful structure for representing object views is to define a set of local image features V along with pairwise relations $E \subset V \times V$ (spatial proximity and (dis)similarity measure) in terms of a weight function $w : E \rightarrow \mathbb{R}_+$, that is an undirected weighted graph $G = (V, E)$ (see Figure 1). Accordingly, a core routine of any recognition system is to compare graphs in different images, or to match graphs against prototypical graphs in some object database.

Finding a good match between two graphs G, G' amounts to compute a permutation of the vertices of one graph so as to become similar to the other one (cf. section 2). It is well-known that this *quadratic assignment problem* is NP-hard [4] and cannot be solved to optimality even for moderately sized problem instances. For the graph shown in Figure 1, for instance, about 10^{44} possible permutations of its vertices exist.

The task of developing good matching algorithms amounts to make a compromise between two contradictory requirements: Firstly, a high-quality solution close to the (unknown) global optimum should be efficiently computable in polynomial time. Secondly, since the matching algorithm is only a particular module of a computer vision system, there should be no need for manual interaction (selection of a starting point or tuning parameters) in order to obtain good performance. It seems to us that this latter algorithmic aspect has received less attention in the computer vision literature than the former one. This motivated us to look for an approach satisfying *both* requirements.

1.2 Related Work

The quadratic assignment problem is a classical problem in the field of combinatorial optimization which is very general in the sense that several well-known problems, for example

the traveling salesman problem, are special cases of it. For a comprehensive review of the quadratic assignment problem we refer to [5].

Classical approaches to the quadratic assignment problem include

- *linearizations* [6, 7] which can be efficiently solved in polynomial time at the cost of a considerably larger number of variables,
- the use of *exact search algorithms*, like the cutting plane method [8] or branch-and-bound [9], which are based on lower bounds [10, 6] of the objective function but can only be applied to small problem instances, or
- *heuristic search algorithms* like tabu search, simulated annealing, or genetic algorithms [11].

A separate class of approaches is formed by *relaxations* of the quadratic assignment problem. Typically, these approaches aim at computing a lower bound of the objective function based on eigenvalue problems [12, 13, 5]. Recently, a convex relaxation approach has been proposed [14] which can be used to compute both a lower bound and the corresponding approximate minimizer. A favorable property due to the convexity of this approach is that the relaxed solution can numerically be computed without the need to select good starting points and tuning parameter values. This will be studied in more detail below.

Another important approach to the quadratic assignment problem is based on the *deterministic annealing* strategy. This approach has been extensively investigated in the neural-network literature [15, 16, 17, 18, 19, 20, 21, 22]. In the context of image segmentation, piecewise-smooth restoration and clustering, deterministic annealing is well-known in the computer vision literature as well [23, 24, 25, 26, 27, 28]. A favorable feature of deterministic annealing is that this strategy can be derived in a theoretically sound way using the maximum entropy principle (e.g., [27]). On the other hand, an obvious disadvantage from the viewpoint of algorithm design are the complex bifurcation phenomena encountered when lowering the annealing parameter [29]. To the best of our knowledge, no control strategy is known which guarantees to reach a “good” local minimum. While excellent experimental results have been reported using deterministic annealing, the dependence of these results on various tuning parameters apparently has not received much attention in the literature. In view of this important aspect we present a thorough comparison of a deterministic annealing approach tailored to the quadratic assignment problem [20, 22] with a convex programming approach [14].

Further important work on the graph-matching problem includes alternative algorithmic approaches like probabilistic relaxation [30], genetic search [31], error-correcting matching [32] or two-step iterative approaches [33], and also specialized work like, for example, simultaneous estimation of transformation geometry [34], or matching trees in terms of the maximum clique of the association graph [35]. A comparison of this variety of approaches is beyond the scope of this paper. We also do not discuss variations of the optimization criterion (see, e.g., [30, 36]), nor do we investigate the task to determine the weight function w in some

optimal way. Rather, by focusing on the optimization problem, we wish to emphasize the advantages of parameter-free convex relaxation in the context of graph-matching and to reveal its performance in comparison to recent approaches which are based on the well-known deterministic annealing framework.

1.3 Contribution

Our work contributes under the aspects *relaxation*, *algorithm design* and *performance evaluation*.

Relaxation has a context-dependent meaning. In this paper, relaxation refers to the strategy of making a combinatorial problem computationally tractable by keeping exactly the objective criterion with respect to feasible combinatorial solutions but weakening the combinatorial constraints. Concerning the quadratic assignment problem, a typical example is to optimize over the orthogonal set of matrices instead of the set of permutation matrices, possibly subject to further constraints like requiring non-negative entries or fixing row and column sums.

Solving (approximately) a combinatorial problem by relaxation has several advantages. Firstly, it is *mathematically* comprehensible where and how approximations are introduced in order to solve the optimization problem. By contrast, many other approaches are based on *algorithmic* modifications, making a theoretical understanding and comparison more difficult. Secondly, relaxed problem formulations compute a lower bound of the objective function, because weakening the constraints gives more degrees of freedom for minimization. As a consequence, different problem relaxations can be *ranked* based on the lower bound they compute: The larger the lower bound, the better the relaxation. Thirdly, relaxations can be used either to directly compute approximate solutions to the original combinatorial problem, or as subroutines in exact branch-and-bound search algorithms. In the latter case the tightness of the lower bound also largely determines the overall performance. Due to these advantages, this paper focuses on relaxations of the quadratic assignment problem.

Algorithm design. A common problem of many approaches concerns the selection of tuning parameters in order to obtain good performance. Concerning the quadratic assignment problem, a representative example – the deterministic annealing strategy – will be examined in more detail below. In this context, *convex* optimization approaches provide an attractive alternative, because the *global* optimum exists under mild conditions and can be computed by established numerical algorithms in polynomial time [37] without *any* additional parameters.

Performance evaluation. A widely-used collection of difficult, real-life benchmark problems exists in the field of combinatorial optimization [38], which has become a standard during the last years. Apparently, this database has not been used in the computer vision literature so far in order to evaluate approaches to the weighted graph matching problem. Besides extensive random ground-truth experiments, our performance evaluation was carried out for problem instances of this database.

1.4 Organisation of the Paper

After stating the problem formally in section 2, we present a hierarchy of relaxations in section 3, the most strongest one being a convex relaxation. For comparison, we sketch two alternative approaches from the literature in section 4, a simple but fast approach based on eigenvalue decomposition, and a more sophisticated deterministic annealing strategy. In order to stress the difference to these non-convex approaches, various aspects of the convex relaxation approach are visualized for a toy example in section 5. In section 6, the results of numerous experiments for both real-life benchmarks from the field of combinatorial optimization and for ground-truth experiments based on computer-generated graphs are summarized. We conclude and indicate further work in section 7.

2 Problem Statement and Definitions

2.1 Notation

We will use the following notation:

X^\top	transpose of the matrix X
I_n	$n \times n$ unit matrix
\mathcal{O}	set of orthogonal matrices X , i.e. $X^\top X = I_n$
\mathcal{E}	set of matrices with unit row and column sums
\mathcal{N}	set of non-negative matrices
Π	set of permutation matrices $\Pi \in \mathcal{O} \cap \mathcal{E} \cap \mathcal{N}$
e	vector of all ones: $e_i = 1, \quad i = 1, \dots, n$
$\text{Tr}[X]$	trace of the matrix X
•	scalar product of two matrices X, Y : $X \bullet Y = \text{Tr}[X^\top Y]$
$\ X\ $	Frobenius norm of the matrix X : $\ X\ = \text{Tr}[X^\top X]^{1/2}$
$\text{vec}[X]$	vector obtained by stacking the columns of the matrix X
\otimes	Kronecker product
$\lambda(X)$	vector of the eigenvalues of the matrix X
$r(X)$	vector of the row sums of the matrix X
$s(X)$	sum of all elements of the matrix X
δ_{ij}	Kronecker delta: $\delta_{ij} = 1$ if $i = j$, and 0 otherwise

2.2 Problem Statement

In this paper, we consider undirected graphs $G = (V, E, w)$ with nodes $V = \{1, \dots, n\}$ and edges $E \subset V \times V$. The weight function $w : E \rightarrow \mathbb{R}_0^+$ typically encodes a similarity measure with respect to pairs of features (i, j) . This measure along with the structure of the graph is represented by the adjacency matrix A_G : $(A_G)_{ij} = w_{ij}$, $i, j = 1, \dots, n$. Since $w_{ij} = w_{ji}$, adjacency matrices are symmetric: $A_G^\top = A_G$.

Let $G = (V_G, E_G, w_G)$ and $H = (V_H, E_H, w_H)$ denote two given graphs. In order to match these two graphs, we want to compute a permutation $\Phi : V_G \mapsto V_H$ of the nodes of G such that the following distance measure is minimized:

$$\sum_{i,j=1}^n (w_{G,\Phi(i)\Phi(j)} - w_{H,ij})^2 \quad (1)$$

Representing the permutation Φ by a permutation matrix $X \in \Pi$, this cost function takes the following form [39] in terms of the adjacency matrices of G and H :

$$f(X) = \|XA_GX^\top - A_H\|^2 \quad (2)$$

For isomorphic graphs exists a permutation matrix such that the minimum value $f(X) = 0$ of the objective function is attained. For features V_G, V_H supplied by an image pre-processing stage, it is unlikely that G and H are isomorphic. In this case we define as the best match the permutation matrix X^* which minimizes f over Π . Thus, the graph matching problem formally reads:

$$f(X^*) = \min_{X \in \Pi} \|XA_GX^\top - A_H\|^2 \quad (3)$$

The minimization problem (3) has a close relationship to the quadratic assignment problem (QAP) in combinatorial mathematics (for a survey, see [5]):

$$\min_{X \in \Pi} \text{Tr}[AXBX^\top + CX^\top] \quad (4)$$

Provided that the graphs have the same number of nodes, this relationship can be seen by reformulating the graph matching objective function as follows:

$$\begin{aligned} f(X) &= \|XA_GX^\top - A_H\|^2 \\ &= \text{Tr}[A_GA_G^\top] + \text{Tr}[A_HA_H^\top] - 2\text{Tr}[A_HXA_G^\top X^\top] \\ &= C_G + C_H - 2\text{Tr}[A_HXA_G^\top X^\top] \end{aligned} \quad (5)$$

Dropping the constant terms $C_G = \text{Tr}[A_GA_G^\top]$ and $C_H = \text{Tr}[A_HA_H^\top]$, we recognize the graph matching problem (3) as a special case of the quadratic assignment problem (4) with $A = A_H$, $B = -A_G$ and $C = 0$. Throughout the remainder of this paper, we can therefore consider the following optimization problem:

$$(QAP) \quad \min_{X \in \Pi} \text{Tr}[AXB^\top X^\top] \quad (6)$$

Remark 1 We note that (1) corresponds to (6) only if $|V_G| = |V_H| = n$. In this paper, we make this simplifying assumption (as did Umeyama [39], for instance) in order to assess the techniques which have been developed for the quadratic assignment problem for the weighted graph matching problem in computer vision. The issue of extending the techniques to *subgraph* matching will be taken up in section 7.

3 Relaxations and Lower Bounds

In this section, we consider various relaxations of problem (6). We will see that a ranking of these approaches can be obtained by virtue of the corresponding lower bounds.

3.1 Orthogonal Relaxation

Relaxing the set Π to $\mathcal{O} \supset \Pi$, Finke et al. [12] suggested the so-called *Eigenvalue Bound (EVB)* which gives a lower bound for the minimization problem (6):

$$(EVB) \quad \min_{X \in \mathcal{O}} \text{Tr}[AXB^T X^T] = \langle \lambda(A), \lambda(B) \rangle_- \quad (7)$$

Here, $\langle \lambda(A), \lambda(B) \rangle_-$ denotes the so-called *minimal scalar product*. This is the scalar product of the vectors $\lambda(A)$ and $\lambda(B)$ containing the eigenvalues of the adjacency matrices A and B ordered as follows: $\lambda_1(A) \leq \lambda_2(A) \leq \dots \leq \lambda_n(A)$ and $\lambda_1(B) \geq \lambda_2(B) \geq \dots \geq \lambda_n(B)$. The matrix X for which the bound (EVB) is attained can be calculated as well. If $U, V \in \mathcal{O}$ diagonalize the adjacency matrices A and B , respectively, and this columns are arranged according to the order of the eigenvalues mentioned above, then $X = UV^T$. It turned out that in many cases this relaxation yields a bound for the minimization problem (6) which is too weak to be useful in practice.

3.2 Projected Eigenvalue Bound

Hadley et al. [13] improved the lower bound (7) by taking into account the constraint set \mathcal{E} in addition to \mathcal{O} . To this end, they parameterized matrices $X \in \mathcal{O} \cap \mathcal{E}$ based on $(n-1) \times (n-1)$ orthogonal matrices $\hat{X} \in \mathcal{O}$ and the relationship

$$X = V\hat{X}V^T + \frac{1}{n}E,$$

where $E = ee^T$ and the $n-1$ columns of the $n \times (n-1)$ matrix V form a basis of the subspace orthogonal to the vector e . Conversely, for any $(n-1) \times (n-1)$ matrix $\hat{X} \in \mathcal{O}$ we have $X = V\hat{X}V^T + \frac{1}{n}E \in \mathcal{O} \cap \mathcal{E}$. The $n \times (n-1)$ matrix V can be calculated using the following scheme:

$$V = \begin{pmatrix} a & a & \dots & a & a \\ 1+b & b & \dots & b & b \\ \vdots & \vdots & \dots & \vdots & \vdots \\ b & b & \dots & b & 1+b \end{pmatrix} \text{ with } a = -\frac{1}{\sqrt{n}} \text{ and } b = -\frac{1}{n + \sqrt{n}}.$$

Using this parameterization, the objective function can be rearranged as follows:

$$\text{Tr}[AXB^T X^T] = \text{Tr}[\hat{A}\hat{X}\hat{B}^T \hat{X}^T] + \text{Tr}[DX] - C_1, \quad (8)$$

where $\hat{A} = V^\top AV$, $\hat{B} = V^\top BV$, $D = \frac{2}{n}r(A)r(B)^\top$ and $C_1 = \frac{1}{n^2}s(A)s(B)$. The authors of [13] suggested to optimize the first two terms on the right hand side of (8) separately, the first one over $\hat{X} \in \mathcal{O}(n-1)$, and the second one over $X \in \Pi$. The latter problem amounts to solve the linear assignment problem

$$\text{LAP}(D) = \min_{X \in \diamond} \text{Tr}[DX], \quad (9)$$

which can be solved using any linear programming solver. As a result, the *Projected Eigenvalue Bound* as a lower bound for the minimization problem (6) is obtained:

$$(PEVB) \quad \langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- + \text{LAP}(D) - C_1 \quad (10)$$

However, a major drawback of this bound is that due to separately minimizing the two terms in (8), a corresponding minimizing matrix X cannot be computed in general. The next subsection shows how one can overcome this drawback by convex relaxation and even achieve a better lower bound than (10).

3.3 Convex Relaxation

Following the work of Anstreicher, Brixius and Wolkowicz [40, 14], we focus on a convex relaxation of the minimization problem (6) in this section. Besides the general arguments discussed in section 1.1, the main motivation for this approach is its ability to compute *both* a tight lower bound and the corresponding matrix X where this bound is attained. In general this is not possible for the bound (10).

As a starting point reconsider the minimization of the first term of the right hand side of equation (8) over the set $\mathcal{O}(n-1)$:

$$\begin{aligned} \min_{\hat{X}} \quad & \text{Tr}[\hat{A}\hat{X}\hat{B}^\top\hat{X}^\top] \\ \text{s.t.} \quad & \hat{X}\hat{X}^\top = I \\ & \hat{X}^\top\hat{X} = I \end{aligned} \quad (11)$$

The Lagrangian dual of this problem reads[40]:

$$\begin{aligned} \max_{\hat{S}, \hat{T}} \quad & \text{Tr}[\hat{S} + \hat{T}] \\ \text{s.t.} \quad & \hat{Q} := (\hat{B} \otimes \hat{A}) - (I \otimes \hat{S}) - (\hat{T} \otimes I) \succeq 0 \\ & \hat{S} = \hat{S}^\top, \hat{T} = \hat{T}^\top \end{aligned} \quad (12)$$

Here $\hat{Q} \succeq 0$ means that \hat{Q} has to be positive semidefinite. The optimal solution for (11), according to (7), is

$$\min_{\hat{X} \in \mathcal{O}(n-1)} \text{Tr}[\hat{A}\hat{X}\hat{B}^\top\hat{X}^\top] = \langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- \quad (13)$$

The duality gap between the optimal solutions of (11) and (12) is zero since interior points exist for both problems (see, e.g., [37]). Hence, the optimal values are the same:

$$\max_{\hat{S}, \hat{T}} \text{Tr}[\hat{S} + \hat{T}] = \langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- \quad (14)$$

The objective function in (11) can be reformulated as follows:

$$\begin{aligned} \text{Tr}[\hat{A}\hat{X}\hat{B}^\top\hat{X}^\top] &= \text{vec}(\hat{X})^\top (\hat{B} \otimes \hat{A}) \text{vec}(\hat{X}) = \text{Tr}[(\hat{B} \otimes \hat{A}) \text{vec}(\hat{X}) \text{vec}(\hat{X})^\top] \\ &= (\hat{B} \otimes \hat{A}) \bullet Y, \end{aligned} \quad (15)$$

where

$$Y = \text{vec}(\hat{X}) \text{vec}(\hat{X})^\top$$

For arbitrary matrices \hat{S} and \hat{T} and $\hat{X} \in \mathcal{O}$ the following equations hold:

$$\begin{aligned} \text{Tr}[\hat{S}] &= \text{Tr}[\hat{S}I] = \text{Tr}[\hat{S}\hat{X}^\top\hat{X}] = \text{Tr}[\hat{X}\hat{S}\hat{X}^\top] = \text{Tr}[I\hat{X}\hat{S}\hat{X}^\top] = (\hat{S} \otimes I) \bullet Y \\ \text{Tr}[\hat{T}] &= \text{Tr}[\hat{T}I] = \text{Tr}[\hat{T}\hat{X}I\hat{X}^\top] = (I \otimes \hat{T}) \bullet Y \end{aligned}$$

Using this, a positive semidefinite form containing \hat{Q} from (12) can be introduced into the objective function $\text{Tr}[\hat{A}\hat{X}\hat{B}^\top\hat{X}^\top]$, if we assume that \hat{S} and \hat{T} are a feasible solution for the dual problem (12):

$$\begin{aligned} \text{Tr}[\hat{A}\hat{X}\hat{B}^\top\hat{X}^\top] &= (\hat{B} \otimes \hat{A}) \bullet Y = (\hat{B} \otimes \hat{A}) \bullet Y + \text{Tr}[\hat{S}] - (\hat{S} \otimes I) \bullet Y + \text{Tr}[\hat{T}] - (I \otimes \hat{T}) \bullet Y \\ &= \text{Tr}[\hat{S} + \hat{T}] + [(\hat{B} \otimes \hat{A}) - (I \otimes \hat{S}) - (\hat{T} \otimes I)] \bullet Y = \text{Tr}[\hat{S} + \hat{T}] + \hat{Q} \bullet Y \\ &= \text{Tr}[\hat{S} + \hat{T}] + \text{vec}(\hat{X})^\top \hat{Q} \text{vec}(\hat{X}) \end{aligned}$$

Choosing \hat{S} and \hat{T} as the optimal solution to (12) we obtain with (14):

$$\text{Tr}[\hat{A}\hat{X}\hat{B}^\top\hat{X}^\top] = \langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- + \text{vec}(\hat{X})^\top \hat{Q} \text{vec}(\hat{X})$$

Finally, substituting this expression as well as all the non-projected variables $\hat{X} = V^\top X V$ etc. into (8), we obtain after an elementary but tedious calculation the Quadratic Programming Bound:

$$(QPB) \quad \text{Tr}[AXB^\top X^\top] = \langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- + \text{vec}(X)^\top Q \text{vec}(X) \quad (16)$$

A comparison with (8) shows that now we have just a single term on the right hand side comprising the unknown matrix X and (16) allows the computation of both a lower bound and the corresponding minimizing matrix X . For the *linear* term in (8), minimizing over the set Π (cf. (9)) is equivalent to minimizing over $\mathcal{E} \cap \mathcal{N}$. Accordingly, Anstreicher and Brixius [14] suggest to minimize the quadratic form in (16) over $\mathcal{E} \cap \mathcal{N}$, i.e. to solve the convex quadratic problem:

$$\begin{aligned} \min \quad & \text{vec}(X)^\top Q \text{vec}(X) \\ \text{s.t.} \quad & Xe = X^\top e = e \\ & X \geq 0 \end{aligned} \quad (17)$$

Here $X \geq 0$ means that all entries of the matrix X have to be non-negative. The following relationship between the bounds (7), (10) and (16) holds[14]:

$$(EVB) \leq (PEVB) \leq (QPB) \leq (QAP) = \min_{X \in \Pi} \text{Tr}[AXB^\top X^\top] \quad (18)$$

Consequently the bound (16) computed by convex programming cannot perform worse than the other bounds. The quality of the corresponding solution X in comparison to other approaches (see next section) will be assessed in section 6.

3.4 Computing a Combinatorial Solution

To obtain a permutation matrix $P \in \Pi$ from the non-integer solution $X \in \mathcal{E} \cap \mathcal{N}$ to (17), a good permutation matrix close to X has to be found. A simple way of doing this is to solve the following linear programming problem:

$$P_0 = \arg \max_{P \in \Pi} \text{Tr}[X^\top P] \quad (19)$$

In this paper, we use a slightly different idea which takes into account that in most cases a linear approximation of the original problem leads to an improvement of the obtained objective function. To this end, we add an unknown matrix Δ to the relaxed solution X so as to give a permutation matrix:

$$P = (X + \Delta) \in \Pi$$

Next, we expand the objective function $\text{Tr}[APB^\top P^\top]$ around X up to linear terms with respect to Δ :

$$\begin{aligned} \text{Tr}[APB^\top P^\top] &= \text{Tr}[A(X + \Delta)B^\top(X + \Delta)^\top] \\ &= \text{Tr}[AXB^\top X^\top] + \text{Tr}[AXB^\top \Delta^\top] + \text{Tr}[A\Delta B^\top X^\top] + \text{Tr}[A\Delta B^\top \Delta^\top] \\ &\approx \text{Tr}[AXB^\top X^\top] + \text{Tr}[AXB^\top \Delta^\top] + \text{Tr}[A\Delta B^\top X^\top] \\ &= -\text{Tr}[AXB^\top X^\top] + \text{Tr}[AXB^\top P^\top] + \text{Tr}[B^\top X^\top AP] \\ &= -\text{Tr}[AXB^\top X^\top] + 2\text{Tr}[B^\top X^\top AP] \end{aligned}$$

As a result, we have to minimize the term $\text{Tr}[B^\top X^\top AP]$ to obtain the combinatorial solution P from the relaxed solution X . This problem can again be solved by linear programming:

$$P_1 = \arg \min_{P \in \Pi} \text{Tr}[B^\top X^\top AP]$$

To see the difference to (19), we put $M = -B^\top X^\top A$ and finally have:

$$P_1 = \arg \max_{P \in \Pi} \text{Tr}[MP] . \quad (20)$$

3.5 The 2opt Post-Processing Heuristics

A simple heuristics called 2opt was proposed in [22] in order to further improve combinatorial solutions computed by more expensive methods. This greedy strategy iteratively exchanges pairs of assignments in the permutation until no further improvement is possible.

4 Non-Convex Approaches

In this section we briefly sketch two approaches that we used for comparison with the convex relaxation approach of section 3.3. The first one was proposed by Umeyama [39] and resembles the spectral relaxation approach of [12]. Furthermore, we consider the deterministic annealing approaches [20] and [22] for which excellent performances are reported in the literature.

4.1 The Approach by Umeyama

Based on the *Eigenvalue Bound* (7), Umeyama [39] proposed the following estimate for the solution of (6):

$$X_{Ume} = \arg \max_{X \in \Pi} \text{Tr}(X^\top |U| |V|^\top) . \quad (21)$$

Here, U and V diagonalize the adjacency matrices A and B , respectively, with the eigenvalues sorted according to (*EVB*), and $|\cdot|$ denotes the matrix consisting of the absolute values taken for each element. (21) is a linear assignment problem which can be efficiently solved by using standard methods like linear programming.

4.2 Graduated Assignment

Gold and Rangarajan [20] and Ishii and Sato [22] independently developed a technique commonly referred to as *graduated assignment* or *soft assign* algorithm. The set of permutation matrices Π is replaced by the convex set $\mathcal{D} = \mathcal{E} \cap \mathcal{N}$ of positive matrices with unit row and column sums (doubly stochastic matrices). In contrast to previous mean-field annealing approaches, the graduated assignment algorithm enforces hard constraints on row *and* column sums, making it usually superior to other deterministic annealing approaches.

The core of the algorithm is an iteration scheme, which computes an approximative solution matrix X at each step of the decreasing annealing schedule. In our description $\beta > 0$ denotes the current annealing parameter; γ is a fixed “self-amplification” parameter, which enforces that the minimum on the set \mathcal{D} is also in Π . Denoting the iteration time step by the superscript, the matrix $X^{(r+1)}$ is calculated as follows (for β fixed):

$$X_{ij}^{(r+1)} = g_i h_j y_{ij}^{(r)} , \quad \text{with } y_{ij}^{(r)} = \exp \left(-\beta \sum_{k,l} (A_{ik} B_{jl} + \delta_{ik} \delta_{jl} \gamma) X_{kl}^{(r)} \right) \quad (22)$$

The scaling coefficients g_i, h_j are computed so that $X^{(r+1)}$ is projected on the set \mathcal{D} using Sinkhorn’s algorithm [20] as inner loop:

$$y_{ij}^{(r,2s)} = \frac{y_{ij}^{(r,2s-1)}}{\sum_k y_{ik}^{(r,2s-1)}} , \quad y_{ij}^{(r,2s+1)} = \frac{y_{ij}^{(r,2s)}}{\sum_k y_{kj}^{(r,2s)}} \quad (23)$$

Stopping criteria based on convergence bounds or the number of iterations have to be established for the inner projection loop and the iteration scheme. For more details, we refer the reader to [20, 22].

Rangarajan et al. [21] showed that this scheme locally converges under mild assumptions. Several studies revealed excellent experimental results. In our experiments, we improved the obtained results with the local 2opt heuristics.

A drawback of the graduated assignment algorithm is that the selection of several “tuning”-parameters is necessary to obtain optimal performance. An annealing schedule has to be set up, which is usually described by three parameters: an initial temperature, the annealing rate, and a final temperature or other stopping criterion [20]. There are theoretically motivated methods that give a lower bound for reasonable initial temperatures based on an analysis of the bifurcation structure of the problem [22]. Nevertheless, careful selection of the parameter greater than this bound can improve the results. The self-amplification parameter also has a lower bound that guarantees the above property that the minimizer of the objective function is in Π . An exhaustive parameter search for the annealing schedule, even below the theoretical bound, may increase the performance. Finally, the stopping criteria also influence the quality of the results. All parameters have in common that their optimal values vary for different problem instances (cf. [22]).

5 Convex Relaxation: An Illustrative Numerical Example

For the purpose of illustration, we apply the convex relaxation approach to a small graph matching problem in this section .

5.1 A Small Graph Matching Problem

In order to graphically visualize the convex relaxation approach, we consider the two small weighted graphs G and H shown in figure 2. Obviously, the best match corresponds to exchanging vertices 2 and 3 in either graph.




Figure 2: Two small sample graphs G and H to be matched

The adjacency matrices of the graphs G and H are:

$$A_G = \begin{pmatrix} 0 & 0.56 & 0.92 \\ 0.56 & 0 & 0.12 \\ 0.92 & 0.12 & 0 \end{pmatrix} \quad A_H = \begin{pmatrix} 0 & 0.99 & 0.22 \\ 0.99 & 0 & 0.02 \\ 0.22 & 0.02 & 0 \end{pmatrix}$$

In this example the objective function of the graph matching problem (3) attains the following values for each of the six possible permutations:

$$1.370, 3.077, 2.01, 3.365, 0.613, 0.261$$

Thus the optimum of this graph matching problem is

$$\text{OPT} = C_G + C_H + 2 \min_{X \in \diamond} \text{Tr}[AXB^\top X^\top] \approx 0.261$$

with $A = A_H, B = -A_G, C_G = \text{Tr}[A_G A_G^\top] \approx 2.349$ and $C_H = \text{Tr}[A_H A_H^\top] \approx 2.058$.

In the following, we visualize permutation matrices X by representing entries $X_{ij} = 1$ graphically by black squares and $X_{ij} = 0$ by white squares. Accordingly, the permutation matrices which lead to the objective function values given above (in the same order) are depicted here:



The last permutation matrix represents an exchange of vertices 2 and 3 and thus corresponds to the global optimum of this graph matching problem.

5.2 Relaxations and Bounds

We calculate the bounds described in section 3 for the problem depicted in figure 2.

Orthogonal Relaxation

The eigenvalue bound leads to the following lower bound:

$$\text{EVB} = C_G + C_H + 2\langle \lambda(A), \lambda(B) \rangle_- \approx 0.023$$

with $C_G = \text{Tr}[A_G A_G^\top] \approx 2.349$, $C_H = \text{Tr}[A_H A_H^\top] \approx 2.058$ and $\langle \lambda(A), \lambda(B) \rangle_- \approx -2.192$. This bound is attained for

$$X \approx \begin{pmatrix} 0.999 & 0.041 & 0.002 \\ -0.014 & 0.316 & 0.948 \\ -0.038 & 0.948 & -0.317 \end{pmatrix}.$$

Due to the small size of our problem, this solution already indicates the optimum despite the obvious weakness of the EVB-bound in general.

Projected Eigenvalue Bound

Using the Projected Eigenvalue Bound, we obtain the following lower bound for our small graph matching problem:

$$\text{PEVB} = C_G + C_H + 2[\langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- + \text{LAP}(D) - C_1] \approx 0.181$$

where $\langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- \approx -0.985$, $\text{LAP}(D) \approx -2.003$, $C_1 \approx -0.875$. Note that this bound is much stronger than the EVB-bound. On the other hand, as mentioned in section 3.2, this approach does not allow to compute a corresponding matrix X for which the PEVB-bound is attained.

Quadratic Programming Bound

The Quadratic Programming Bound gives:

$$\text{QPB} = C_G + C_H + 2[\langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- + \min_{X \in \mathcal{E} \cap \mathcal{N}} \text{vec}(X)^\top Q \text{vec}(X)] \approx 0.215$$

Here the minimization of the quadratic term results in $\min_{X \in \mathcal{E} \cap \mathcal{N}} \text{vec}(X)^\top Q \text{vec}(X) \approx -1.111$ and the bound is attained for

$$X \approx \begin{pmatrix} 0.747 & 0.000 & 0.253 \\ 0.253 & 0.000 & 0.747 \\ 0.000 & 1.000 & 0.000 \end{pmatrix}.$$

As predicted, this bound is superior to the PEVB-bound. So summarizing, for the numerical example considered here the ranking (18) of these bounds reads:

$$\text{EVB} \approx 0.023 \leq \text{PEVB} \approx 0.181 \leq \text{QPB} \approx 0.215 \leq \text{OPT} \approx 0.261 \quad (24)$$

5.3 Visualization

To illustrate how the convex relaxation approximates the original combinatorial problem, we inspect graphically the original objective function

$$f_{\text{orig}}(X) = C_G + C_H + 2\text{Tr}[AXB^\top X^\top]$$

along with its convex relaxation

$$f_{\text{convex}}(X) = C_G + C_H + 2[\langle \lambda(\hat{A}), \lambda(\hat{B}) \rangle_- + \text{vec}(X)^\top Q \text{vec}(X)]$$

for a few one-dimensional paths $X(\alpha)$ through the relaxed solution set defined by $X \in \mathcal{E} \cap \mathcal{N}$. It is well known (Birkhoff-von Neumann theorem) that this set is just the convex hull of the original feasible set, i.e. the permutation matrices $X \in \Pi$. Hence, all paths

$$X(\alpha) = \alpha X_2 + (1 - \alpha) X_1, \quad \alpha \in [0, 1]$$




Figure 3: The original objective function and its convex relaxation along paths through the relaxed solution set.

between extreme points $X_1, X_2 \in \Pi$ go through the interior relaxed solution set, and we can graphically explore the two cost functions above by plotting their graphs over various paths. Figures 3 and 4 show several paths and illustrate the following facts:

- At the end–points of all paths, the two cost functions coincide because the relaxed approach does not change the original objective function at the original feasible set (cf. Section 1.3).
- The original objective function is non convex in the relaxed solution set and thus exhibits local minima. This is not the case for the objective function of the convex relaxation.
- The plot on the right hand side of Figure 4 illustrates how the lower bound is attained by the convex relaxation approach. Furthermore, the point where this bound is attained is close to the global optimum (the end–point on the right) due to the tightness of the lower bound.

In summary, these Figures illustrate that the convex relaxation “feels” where the “good” minimum lies, and that a nearby point can be computed without any initialisation problem or parameter tuning!

6 Experiments and Discussion

This section has two parts. In the first part we investigate the performance of the convex optimization relaxation. To this end, we compare the corresponding lower bounds with the combinatorial solutions of several benchmark problems from the QAPLIB–collection [38]. The QAPLIB is a public library of very difficult real–life quadratic assignment problems

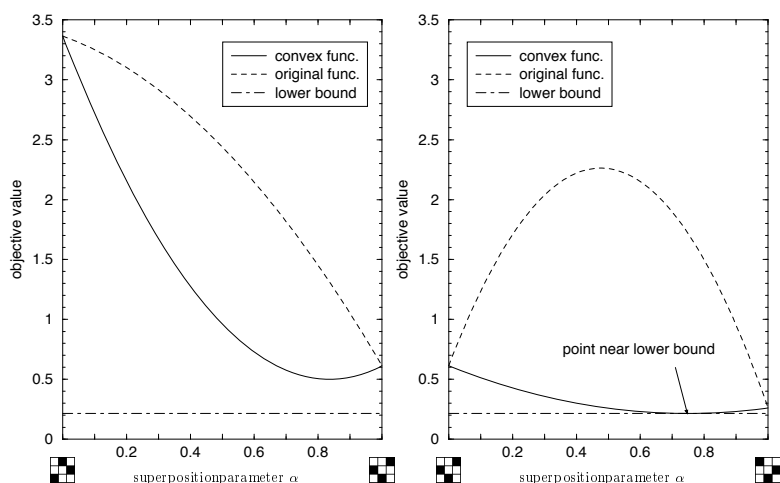


Figure 4: The original objective function and its convex relaxation along paths through the relaxed solution set. On the right, the plot illustrates how the lower bound is attained at a point close to the global optimum (end–point on the right).

which can be used to evaluate and to compare the performance of any quadratic assignment approach. In the second part we present statistical results computed for a large set of randomly generated graphs (including ground–truth).

Abbreviations

The following abbreviations are used within the tables of this section. f represents in all cases the value calculated by the objective function (cf.(6))

$$f(X) = \text{Tr}[AXB^\top X^\top],$$

with an $X \in \Pi$. The subscript of f shows how and with which approach the $X \in \Pi$ was obtained.

f^* : value of the objective function (6) at the global optimum $X^* \in \Pi$

EVB : the eigenvalue bound (7)

$PEVB$: the projected eigenvalue bound (10)

QPB : the quadratic programming bound (16)

f_{QPB} : value of the objective function in (6) using the permutation matrix obtained with (19) from the QPB solution.

f_{QPB}^1 : value of the objective function in (6) using the permutation matrix obtained with (20) from the QPB solution.

f_{GA} : value of the objective function in (6) using the permutation matrix obtained by the graduated assignment algorithm (22).

f_{Ume} : value of the objective function (6) using the permutation matrix obtained by the approach from Umeyama (21)

An additional “+”-sign (e.g. f_{QPB+} , f_{QPB+}^1 , f_{GA+} , f_{Ume+}) indicates that the 2opt-heuristics was used as a post-processing step to further improve the permutation matrix found.

6.1 QAPLIB Benchmark Experiments

Quality of the Relaxations

The quality of the various relaxation approaches, namely the eigenvalue bound (EVB), the projected eigenvalue bound ($PEVB$) and the quadratic programming bound (QPB), can be assessed by measuring how close these bounds are to the global optimum (see (18)).

Problem	f^*	EVB	$PEVB$	QPB
chr12c	11156	-127514	-24375	-22648
chr15a	9896	-190769	-52468	-48539
chr15c	9504	-186403	-50295	-47409
chr20b	2298	-30995	-8051	-7728
chr22b	6194	-66432	-22126	-20995
esc16b	292	-230	250	250
rou12	235528	-274122	200024	205461
rou15	354210	-424419	296705	303487
rou20	725522	-739730	597045	607362
tai10a	135028	-181950	112528	116260
tai12a	224416	-284261	193124	199378
tai15a	388214	-414351	325019	330205
tai17a	491812	-496403	408910	415578
tai20a	703482	-714901	575831	584942
tai30a	1818146	-1505553	1500406	1517829
tai35a	2422002	-2015233	1941622	1958998
tai40a	3139370	-2559063	2484371	2506806

Table 1: Bounds computed for QAPLIB-problems.

Table 1 shows the results for problems drawn from the QAPLIB [38]. The first column comprises labels indicating the problem and the number $|V|$ of vertices of a data set from the QAPLIB. The second column shows the value of the objective function at the global optimum. The corresponding lower bounds computed by the relaxation approaches are listed in the remaining columns.

Since zero is a trivial lower bound, a negative sign indicates that the relaxation is not tight. This happens for most problems with the EV-bound which therefore can be considered not to be useful, and for some problems with the other bounds as well (hence these problems seem to be most difficult).

Furthermore, table 1 confirms the relationship (18), and that the convex relaxation approach gives the best lower bound.

Comparison to Spectral Decomposition and Graduated Assignment

We compare the combinatorial solutions obtained with the convex relaxation approach with those computed with the graduated assignment approach [20, 22] and the spectral decomposition approach by Umeyama [39]. Table 2 shows the results in the same way as table 1, but

Problem	f^*	f_{QPB}	f_{QPB+}	f_{QPB}^{\downarrow}	f_{QPB+}^{\downarrow}	f_{GA}	f_{GA+}	f_{Ume}	f_{Ume+}
chr12c	11156	20306	15860	27912	13088	19014	11186	40370	11798
chr15a	9896	26132	14454	20640	13540	30370	11062	60986	17390
chr15c	9504	29862	17342	19436	12754	23686	13342	76318	13338
chr20b	2298	6674	2858	7276	3832	6290	2650	10022	3294
chr22b	6194	9942	6848	8958	6902	9658	6732	13118	7418
esc16b	292	296	292	312	292	298	292	306	292
rou12	235528	278834	246712	266864	241802	273438	246282	295752	251848
rou15	354210	381016	371480	394192	374000	457908	359748	480352	384018
rou20	725522	804676	746636	795578	757270	840120	738618	905246	765872
tai10a	135028	165364	143260	154282	139524	168096	135828	189852	147838
tai12a	224416	263978	237200	246424	238902	263778	224416	294320	252044
tai15a	388214	455778	399732	432610	390782	451164	400328	483596	405442
tai17a	491812	550852	513170	545410	526518	589814	505856	620964	526814
tai20a	703482	799790	740696	752896	726038	871480	724188	915144	775456
tai30a	1818146	1996442	1883810	1979530	1872722	2077958	1886790	2213846	1875680
tai35a	2422002	2720986	2527684	2677688	2511800	2803456	2496524	2925390	2544536
tai40a	3139370	3529402	3243018	3411278	3277450	3668044	3249924	3727478	3282284

Table 2: Results of the QAPLIB benchmark experiments (see text).

now only combinatorial solutions are shown. A “+”-sign indicates that the 2opt-heuristics was used as a post-processing step to improve the solution. The difference between f_{QPB} and f_{QPB}^{\downarrow} is that linearization was used to “round” the relaxed convex programming solution to a combinatorial solution in the latter case (see section 3.4).

The columns labeled with f_{GA} and f_{Ume} show the results obtained for the graduated assignment approach [20, 22] and for the approach by Umeyama [39]. It should be noted that considerable care was taken to find out optimal parameter values for the graduated assignment approach for each data set [42].

The following conclusions can be drawn from the results shown in table 2:

- The convex relaxation approach f_{QPB} and the soft-assign approach f_{GA} have similarly good performance, despite the fact that the latter approach is much more intricate from the optimization point-of-view and involves a couple of tuning parameters which have to be (and were) optimized by hand.
- The approach of Umeyama f_{Ume} based on spectral decomposition is not as competitive.
- Using the simple 2opt greedy-strategy as a post-processing step significantly improves the solution in most cases.

6.2 Random Ground-Truth Experiments

In this subsection we discuss our results obtained for two different ground-truth experiments. In the first experiment we created many problem instances (6) by independently computing two different random graphs with the same number of vertices. In the second experiment we computed a large collection of random graphs along with slightly perturbed and randomly permuted “copies” of these graphs.

Random Graphs

In this experiment we created many problem instances (6) by independently computing two different random graphs with the same number of vertices. The probability that an edge is present in the underlying complete graph was about 0.3. Figure 5 shows an example in order to visualize the edge-density of such graphs. The global optimum for (6) was computed using an exact search algorithm. The global optimum was used to calculate the ratio of the suboptimal objective value to the best objective value for each problem instance. Table 3 summarizes our results based on this ratio. It shows the statistics (mean, worst case and the best case) for three experiments with different sizes of the graphs ($n = 9, 11, 15$). The number of problem instances for each experiment is shown in angular brackets. The number of correctly found matchings without/with the 2opt heuristics as post-processing step are shown in round brackets. The following conclusions can be drawn from the results shown in table 3:

- The soft-assign approach performs somewhat better for these experiments than the convex relaxation approach but the latter needs no tuning parameters which have to be optimized by hand.
- With increasing problem size the performance decreases for all three approaches.
- The approach of Umeyama f_{Ume} based on spectral decomposition is not as competitive.
- Using the simple 2opt greedy-strategy as a post-processing step significantly improves the results.

	f_{QP}^1/f^*			f_{Ume}/f^*			f_{GA}/f^*		
	mean	worst case	best case	mean	worst case	best case	mean	worst case	best case
$n=9$ [128]	(22/55)			(7/29)			(31/55)		
2opt	0.88765	0.43810	1	0.638244	0.065173	1	.948342	.7756129	1
	0.97130	0.79256	1	0.928304	0.753007	1	.969914	.843046	1
$n=11$ [42]	(3/10)			(0/7)			(7/10)		
2opt	0.83043	0.56268	1	0.636159	0.295194	0.998591	.940740	.8338586	1
	0.95760	0.85043	1	0.933206	0.811326	1	.958863	.8434407	1
$n=15$ [99]	(0/2)			(0/1)			(4/11)		
2opt	0.78726	0.52307	0.938917	0.225983	0.131333	0.863508	.916225	.105164	1
	0.92195	0.77956	1	0.890131	0.74688	1	.95763	.820596	1

Table 3: Statistics of the results of random ground-truth experiments (see text).

Perturbed Graphs

In the second series of experiments we computed a large collection of random graphs along with slightly perturbed and randomly permuted “copies” of these graphs. The weights of the second graph were perturbed by a normally distributed factor with standard deviation $\sigma = 0.1$ around 1. The results for this kind of experiments are shown in table 4 which has the same structure as table 3. For larger problems (more than 15 vertices) where computing the global




Figure 5: A randomly generated graph with 15 vertices and a probability of about 0.3 for the presence of an edge.

optimum was too expensive, we assumed the optimal permutation to be the inverse of the random permutation matrix which was used to compute the second graph of each pair. In some cases this was not true and hence a different permutation with a lower objective value could be found by the algorithms. This explains why some of the quotients in table 4 have a value greater than 1.

In summary, the statistics of our results shown in table 4 reveal that in almost every case of these “low-level noise” experiments the optimal permutation was found by the quadratic programming approach.

	f_{QPB}^1/f^*			f_{Ume}/f^*			f_{GA}/f^*		
	mean	worst case	best case	mean	worst case	best case	mean	worst case	best case
n=9 [155]	(154/155)			(142/154)			(144/151)		
2opt	0.999996	0.999382	1	0.986481	0.463282	1	.997206	.859380	1
	1	1	1	0.999883	0.981862	1	.998154	.859380	1
n=15 [183]	(183/183)			(163/175)			(176/181)		
2opt	1	1	1	0.974078	0.379189	1	.998347	.833787	1
	1	1	1	0.993836	0.718871	1	.998484	.833787	1
n=20 [173]	(173/173)			(148/163)			(167/171)		
2opt	1	1	1	0.977225	0.475711	1	.998205	.855257	1
	1	1	1	0.991662	0.772512	1	.998338	.855257	1
n=25 [169]	(169/169)			(64/123)			(126/143)		
2opt	1.00001	1	1.00155	0.848105	0.216079	1	.966097	.491432	1.001550
	1.00002	1	1.00155	0.960519	0.602629	1.00099	.9748815	.686842	1.001550

Table 4: Statistics of the results of perturbed graph experiments (see text).

7 Conclusion and Further Work

We showed that the convex programming approach is a competitive approach for finding suboptimal solutions to the weighted graph-matching problem. We compared the convex approach with both a recent deterministic annealing approach and an approach based on the eigenvalue decomposition. The performance of the latter approach is worse whereas the deterministic annealing approach performs similarly or slightly better, but uses parameters values which were optimized by hand. The advantage of the convex approach is that no “tuning” parameters have to be determined at all. Furthermore, in contrast with the deterministic annealing approach, the convex approach provides a lower bound and thus can be used as a subroutine within an exact search strategy like branch-and-bound. Our results show that it is an attractive direction of research for solving relational matching problems in the context of view-based object recognition.

Our further work will focus on the case of graphs with an unequal number of vertices: $|V_G| \neq |V_H|$. If this difference is small, our approach can be applied by filling up the smaller graph with “virtual nodes”. In general, of course, this is not a satisfying way. The consequence of different numbers of vertices is that the unknown permutation matrix X becomes a matching matrix, and that either of the two constants C_G, C_H in the objective function (5) changes to a term which depends on X , too. In our further work we will extend the approach presented here to this more general case.

Acknowledgement

The authors are grateful to Jens Keuchel and Daniel Cremers for scrutinizing the manuscript and many suggestions.

References

- [1] M. Herbert, J. Ponce, T. Boult, and A. Gross, editors. *Object Representation in Computer Vision*, volume 994 of *Lect. Not. Comp. Sci.* Springer-Verlag, 1995.
- [2] H.H. Bülthoff and S. Edelman. Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proc. Nat. Acad. Science.*, 92:60–64, 1992.
- [3] W. Förstner. A framework for low level feature extraction. In J.O. Eklundh, editor, *Computer Vision - ECCV '94*, volume 801 of *Lect. Not. Comp. Sci.*, pages 61–70. Springer-Verlag, 1994.
- [4] S. Sahni and T. Gonzalez. P-complete approximation problems. *J. ACM*, 23:555–565, 1976.

- [5] R.E. Burkard, E. Cela, P.M. Pardalos, and L.S Pitsoulis. The quadratic assignment problem. In P.M. Pardalos and D.-Z. Du, editors, *Handbook of Combinatorial Optimization*, pages 241–338. Kluwer Acad. Publishers, 1998.
- [6] E.L. Lawler. The quadratic assignment problem. *Management Science*, 9:586–599, 1963.
- [7] H.A. Almohamad and S.O. Duffuaa. A linear programming approach for the weighted graph matching problem. *IEEE Trans. Patt. Anal. Mach. Intell.*, 15(5):522–525, 1993.
- [8] M.S. Bazaraa and H.D. Sherali. On the use of exact and heuristic cutting plane methods for the quadratic assignment problem. *J. Operations Res. Soc.*, 33:991–1003, 1982.
- [9] P.M. Pardalos, K.G. Ramakrishnan, M.G.C. Resende, and Y. Li. Implementation of a variable reduction based lower bound in a branch and bound algorithm for the quadratic assignment problem. *SIAM J. Optimiz.*, 7:280–294, 1997.
- [10] P. Gilmore. Optimal and suboptimal algorithms for the quadratic assignment problem. *SIAM J. Applied Math.*, 10:305–313, 1962.
- [11] E. Aarts and J.K. Lenstra, editors. *Local Search in Combinatorial Optimization*, Chichester, 1997. Wiley & Sons.
- [12] G. Finke, R.E. Burkard, and F. Rendl. Quadratic assignment problems. *Annals of Discrete Mathematics*, 31:61–82, 1987.
- [13] S.W. Hadley, F. Rendl, and H. Wolkowicz. A new lower bound via projection for the quadratic assignment problem. *Math. of Operations Research*, 17:727–739, 1992.
- [14] N. Brixius and K. Anstreicher. Solving quadratic assignment problems using convex quadratic programming relaxations. *Optimization Methods and Software*, 16(1–4):49–68, 2001.
- [15] R. Kree and A. Zippelius. Recognition of topological features of graphs and images in neural networks. *J. Phys. A*, 21:L813–L818, 1988.
- [16] C. Peterson and B. Söderberg. A new method for mapping optimization problems onto neural networks. *Int. J. of Neural Systems*, 1:3–22, 1989.
- [17] A.L. Yuille. Generalized deformable models, statistical physics, and matching problems. *Neural Comp.*, 2:1–24, 1990.
- [18] P. Simic. Constrained nets for graph matching and other quadratic assignment problems. *Neural Computation*, 3(2):268–281, 1991.

- [19] J.J. Kosowsky and A.L. Yuille. The invisible hand algorithm: Solving the assignment problem with statistical physics. *Neural Networks*, 7(3):477–490, 1994.
- [20] S. Gold and A. Rangarajan. A graduated assignment algorithm for graph matching. *IEEE Trans. Patt. Anal. Mach. Intell.*, 18(4):377–388, 1996.
- [21] A. Rangarajan, A. Yuille, and E. Mjolsness. Convergence properties of the softassign quadratic assignment algorithm. *Neural Computation*, 11(6):1455–1474, 1999.
- [22] S. Ishii and M. Sato. Doubly constrained network for combinatorial optimization. *Neurocomputing*, 43, 239-257 (2002).
- [23] A. Blake and A. Zisserman. *Visual Reconstruction*. MIT Press, 1987.
- [24] Y.G. Leclerc. Constructing simple stable descriptions for image partitioning. *Int. J. of Comp. Vision*, 3(1):73–102, 1989.
- [25] D. Geiger and F. Girosi. Parallel and deterministic algorithms from mrf's: Surface reconstruction. *IEEE Trans. Patt. Anal. Mach. Intell.*, 13(5):401–412, 1991.
- [26] D. Geiger and A. Yuille. A common framework for image segmentation. *Int. J. of Comp. Vision*, 6(3):227–243, 1991.
- [27] T. Hofmann and J. Buhmann. Pairwise data clustering by deterministic annealing. *IEEE Trans. Patt. Anal. Mach. Intell.*, 19(1):1–14, 1997.
- [28] A.V. Rao, D.J. Miller, K. Rose, and A. Gersho. A deterministic annealing approach for parsimonious design of piecewise regression models. *IEEE Trans. Patt. Anal. Mach. Intell.*, 21(2):159–173, 1999.
- [29] M. Sato and S. Ishii. Bifurcations in mean-field-theory annealing. *Physical Review E*, 53(5):5153–5168, 1996.
- [30] W.J. Christmas, J. Kittler, and M. Petrou. Structural matching in computer vision using probabilistic relaxation. *IEEE Trans. Patt. Anal. Mach. Intell.*, 17(8):749–764, 1995.
- [31] A.D.J. Cross, R.C. Wilson, and E.R. Hancock. Inexact graph matching using genetic search. *Pattern Recog.*, 30(6):953–970, 1997.
- [32] B.T. Messmer and H. Bunke. A new algorithm for error-tolerant subgraph isomorphism detection. *IEEE Trans. Patt. Anal. Mach. Intell.*, 20(5):493–504, 1998.
- [33] B. Luo and E.R. Hancock. Structural graph matching using the em algorithm and singular value decomposition. *IEEE Trans. Patt. Anal. Mach. Intell.*, 23(10):1120–1136, 2001.

- [34] A.D.J. Cross and E.R. Hancock. Graph-matching with a dual-step em algorithm. *IEEE Trans. Patt. Anal. Mach. Intell.*, 20(11):1236–1253, 1998.
- [35] M. Pelillo, K. Siddiqi, and S.W. Zucker. Matching hierarchical structures using association graphs. *IEEE Trans. Patt. Anal. Mach. Intell.*, 21(11):1105–1120, 1999.
- [36] H. Bunke. Error correcting graph matching: On the influence of the underlying cost function. *IEEE Trans. Patt. Anal. Mach. Intell.*, 21(9):917–922, 1999.
- [37] Y. Nesterov and A. Nemirovskii. *Interior Point Polynomial Methods in Convex Programming*. SIAM, 1994.
- [38] R.E. Burkard, S. Karisch, and F. Rendl. Qaplib – a quadratic assignment problem library. *J. Global Optimization*, 10:391–403, 1997.
- [39] S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE Trans. Patt. Anal. Mach. Intell.*, 10(5):695–703, 1988.
- [40] K. Anstreicher and H. Wolkowicz. On langrangian relaxation of quadratic matrix constraints. *SIAM J. Matrix Anal. Appl.*, 22(1):41-55, 2000
- [41] K.M. Anstreicher and N.W. Brixius. A new bound for the quadratic assignment problem based on convex quadratic programming. Technical report, Dept. of Management Sciences, University of Iowa, 1999.
- [42] S. Roth. Analysis of a deterministic annealing method for graph matching and quadratic assignment problems in computer vision. Master’s thesis, CVGPR–group, University of Mannheim, May 2001.