Evaluation of Convex Optimization Techniques for the Weighted Graph-Matching Problem in Computer Vision

Christian Schellewald, Stefan Roth, Christoph Schnörr

Computer Vision, Graphics, and Pattern Recognition Group
Dept. Mathematics and Computer Science
University of Mannheim, D-68131 Mannheim, Germany
{cschelle,roths,schnoerr}@ti.uni-mannheim.de

Abstract. We present a novel approach to the weighted graph-matching problem in computer vision, based on a convex relaxation of the underlying combinatorial optimization problem. The approach always computes a lower bound of the objective function, which is a favorable property in the context of exact search algorithms. Furthermore, no tuning parameters have to be selected by the user, due to the convexity of the relaxed problem formulation.
For comparison, we implemented a recently published deterministic annealing approach and conducted numerous experiments for both established benchmark experiments from combinatorial mathematics, and for random ground-truth experiments using computer-generated graphs. Our results show similar performance for both approaches. In contrast to the convex approach, however, four parameters have to be determined by hand for the annealing algorithm to become competitive.

1 Introduction

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 provide evidence for *view-based* object representations [2] in human vision.

A common and powerful representation format for object views is a set of local image features V along with pairwise relations E (spatial proximity and (dis)similarity measure), that is an undirected graph $G = (V, E)$. In this paper, we will discuss the application of a novel convex optimization technique to the problem of matching relational representations of object views.

Relations to previous work. There are numerous approaches to graph-matching in the literature (e.g., [3–8]). Our work differs from them with respect to the following points:

1. We focus on *problem relaxations*, i.e. the optimization criterion equals the original one but is subject to weaker constraints. As a consequence, such approaches compute a *lower bound* of the original objective function which has to be minimized.
2. The *global* optimum of the *relaxed* problem can be computed with polynomial-time complexity.

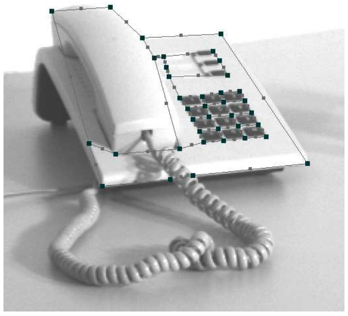

Fig.1 A graph based on features described in [9] using corresponding public software (http://www.ipb.uni-bonn.de/ipb/projects/fex/fex.html). This graph has $|V| = 38$ nodes.

The first property above is necessary for combining the approach with an exact search algorithm where lower bounds of the original objective function are needed. Furthermore, it allows to compare different approaches by simply ranking the corresponding lower bounds.

The second property is important since graph-matching belongs to the class of NP-hard combinatorial problems. Matching two graphs with, say, $|V| = 20$ nodes gives $\sim 10^{18}$ possible matches. Typical problem instances however (see Fig. 1) comprise $|V| > 20$ nodes and thus motivate to look for tight problem relaxations to compute good suboptimal solutions in polynomial time.

Contribution. We discuss the application of novel convex optimization techniques to the graph-matching problem in computer vision.
First, we sketch a recently published deterministic annealing approach [6, 10] which stimulated considerable interest in the literature due to its excellent performance in numerical experiments. Unfortunately, this approach cannot be interpreted as a relaxation of the graph-matching problem and requires the selection of (at least) four parameters to obtain optimal performance (Section 3).
Next we consider the relaxed problems proposed in [11, 3] and show that, by using convex optimization techniques based on the work [12], a relaxation of the graph-matching problem is obtained with equal or better performance than the other approaches (Section 4). Moreover, due to convexity, no parameter selection is required.
In Section 5, we report extensive numerical results with respect to both benchmark-experiments [13] from the field of combinatorial mathematics, and random ground-truth experiments using computer-generated graphs.

Remark. Note that, in this paper, we are exclusively concerned with the *optimization procedure* of the graph-matching problem. For issues related to the *design of the optimization criterion* we refer to, e.g., [4, 14].

Notation.

X^t:	transpose of a matrix X
$\mathcal{O}$:	set of orthogonal $n \times n$-matrices X, i.e. $X^t X = I$ (I: unit matrix)
$\mathcal{E}$:	matrices with unit row and column sums
$\mathcal{N}$:	set of non-negative matrices
Π:	set of permutation matrices $X \in \mathcal{O} \cap \mathcal{E} \cap \mathcal{N}$
e:	one-vector $e_i = 1, \quad i = 1, \ldots, n$
vec $[A]$:	vector obtained by stacking the columns of some matrix A
$\lambda(A)$:	vector of eigenvalues of some matrix A

2 Problem statement

Let $G = (V, E)$, $G' = (V', E')$ denote two weighted undirected graphs with $|V| = |V'| = n$, weights $\{w_{ij}\}, \{w'_{ij}\}$, and adjacency matrices $A_{ij} = w_{ij}$, $A'_{ij} = w'_{ij}$, $i, j = 1, \ldots, n$. Furthermore, let ϕ denote a permutation of the set $\{1, \ldots, n\}$ and $X \in \Pi$ the corresponding permutation matrix, that is $X_{ij} = 1$ if $\phi(i) = j$ and $X_{ij} = 0$ otherwise. The weight functions $w, w' : E \subset V \times V \to \mathbb{R}_0^+$ encode (dis)similarity measures of local image features V_i, $i = 1, \ldots, n$, which we assume to be given in this paper. We are interested in matching graphs G and G' by choosing a permutation ϕ^* such that

$$\phi^* = \arg\min_{\phi} \sum_{i,j} (w_{\phi(i)\phi(j)} - w'_{ij})^2 \ .$$

By expanding and dropping constant terms, we obtain the equivalent problem:

$$\min_{\phi}(-\sum_{i,j} w_{\phi(i)\phi(j)} w'_{ij}) = \min_{X}(-\mathrm{tr}(A' X A X^t)) \ ,$$

with $\mathrm{tr}(\cdot)$ denoting the trace of a matrix. Absorbing the minus sign, we arrive at the following *Quadratic Assignment Problem (QAP)* with some arbitrary, symmetric matrices A, B:

$$(QAP) \qquad \min_{X \in \Pi} \mathrm{tr}(A X B X^t) \ . \tag{1}$$

3 Graduated assignment

Gold and Rangarajan [6] and Ishii and Sato [10] 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 the following iteration scheme, where $\beta > 0$ denotes the annealing parameter and the superscript denotes the iteration time step (for β fixed):

$$X_{ij}^{(r+1)} = g_i h_j y_{ij}^{(r)} \ , \quad \text{with } y_{ij}^{(r)} = \exp\left(-\beta \sum_{k,l} A_{ik} B_{jl} X_{kl}^{(r)}\right) \tag{2}$$

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 [6] as inner loop.

This scheme locally converges under mild assumptions [15]. Several studies revealed excellent experimental results. In our experiments, we improved the obtained results with a local 2opt heuristics which iteratively improve the objective function by exchanging two rows and columns of the found permutation matrix until no improvement in the objective function is possible, as proposed in [10].

A drawback of this approach is that the selection of several "tuning"-parameters is necessary to obtain optimal performance, namely:

- the parameter β related to the annealing schedule,
- a "self-amplification" parameter enforcing integer values, and
- two stopping criteria with respect to the two iteration loops in (2).

Furthermore, the optimal parameter values vary for different problem instances (cf. [10]). For more details, we refer to [16].

4 Convex Approximations

In this section, we discuss a *convex* approximation to the weighted graph-matching problem (1). For more details and proofs, we refer to [16].
As explained in Section 1, our motivation is twofold: Firstly, the need to select parameter values (cf. previous section) is quite inconvenient when using a graph-matching approach as a part within a computer vision system. Convex optimization problems admit algorithmic solutions without any further parameters. Secondly, we focus on problem relaxations providing lower bounds of the objective criterion (1), which then can be used in the context of exact search algorithms.

4.1 Orthogonal relaxation and eigenvalue bounds

Replacing the set Π by $\mathcal{O} \supset \Pi$, Finke et al. [11] proved the following so called *Eigenvalue Bound* (EVB) as a lower lower bound of (1):

$$(EVB) \qquad \min_{X \in \mathcal{O}} \operatorname{tr}(AXBX^t) = \big(\lambda(A)\big)^t \lambda(B) \ , \tag{3}$$

with $\lambda(A), \lambda(B)$ sorted such that $\lambda_1(A) \geq \cdots \geq \lambda_n(A)$ and $\lambda_1(B) \leq \cdots \leq \lambda_n(B)$. This bound can be improved to give the *Projected Eigenvalue Bound* ($PEVB$) by further constraining the set of admissible matrices [17], but in contrast to the approach sketched in Section 4.3 this does not produce a matrix X for which the bound ($PEVB$) is attained.

4.2 The approach by Umeyama

Based on (3), Umeyama [3] proposed the following estimate for the solution of (1):

$$\hat{X}_{Ume} = \arg\max_{X \in \Pi} \operatorname{tr}(X^t \, |U| \, |V|^t) \, . \tag{4}$$

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 value taken for each element. (4) is a linear assignment problem which can be efficiently solved by using standard methods like linear programming.

4.3 Convex relaxation

Anstreicher and Brixius [12] improved the projected eigenvalue bound $(PEVB)$ introduced in Section 4.1 to the *Quadratic Programming Bound* (QPB):

$$(QPB) \qquad \left(\lambda(\tilde{A})\right)^t \lambda(\tilde{B}) + \min_{X \in \mathcal{E} \cap \mathcal{N}} \operatorname{vec}\,[X]^t Q \operatorname{vec}\,[X] \, , \tag{5}$$

where $\tilde{A} = P^t A P$, $\tilde{B} = P^t B P$, with P being the orthogonal projection onto the complement of the 1D-subspace spanned by the vector e, and where the matrix Q is computed as solution to the Lagrangian dual problem of the minimization problem (3) (see [12, 16] for more details). Notice that both the computation of Q and minimizing (QPB) are *convex* optimization problems. Let $\tilde{X}$ denote the global minimizer of (5). Then we compute a suboptimal solution to (1) by solving the following linear assignment problem:

$$\hat{X}_{QPB} = \arg\min_{X \in \Pi} \operatorname{tr}(X^t \tilde{X}). \tag{6}$$

The bounds presented so far can be ranked as follows:

$$(EVB) \leq (PEVB) \leq (QPB) \leq (QAP) = \min_{X \in \Pi} \operatorname{tr}(AXBX^t) \, . \tag{7}$$

We therefore expect to obtain better solutions to (1) using (6) than using (4). This will be confirmed in the following section.

5 Experiments

We conducted extensive numerical experiments in order to compare the approaches sketched in Sections 3 and 4. The results are summarized in the following. Two classes of experiments were carried out:

– We used the QAPLIB-library [13] from combinatorial mathematics which is a collection of problems of the form (1) which are known to be "particularly difficult".

– Furthermore, we used large sets of computer-generated random graphs with sizes up to $|V| = 15$ such that (i) the global optimum could be computed as ground-truth by using an exact search algorithm, and (ii) significant statistical results could be obtained with respect to the quality of the various approaches.

QAPLIB benchmark experiments.
Table 1 shows the results computed for several QAPLIB-problems. The following abbreviations are used:

QAP: name of the problem instance (1) taken from the library
X^*: value of the objective function (1) at the global optimum
QPB: the quadratic programming bound (5)
$\hat{X}_{QPB}$: value of the objective function (1) using $\hat{X}_{QPB}$ from (6)
$\hat{X}_{QPB}+$: $\hat{X}_{QPB}$ followed by the 2opt greedy-strategy
$\hat{X}_{GA}$: value of the objective function (1) using $\hat{X}$ from (2)
$\hat{X}_{GA}+$: $\hat{X}_{GA}$ followed by the 2opt greedy-strategy
$\hat{X}_{Ume}$: value of the objective function (1) using (4)
$\hat{X}_{Ume}+$: $\hat{X}_{Ume}$ followed by the 2opt greedy-strategy

The 2opt greedy-strategy amounts to iteratively exchanging two rows and columns of the matrix $\hat{X}$ as long as an improvement of the objective function is possible [10].

By inspection of table 1, three conclusions can be drawn:

– The convex relaxation approach $\hat{X}_{QPB}$ and the soft-assign approach $\hat{X}_{GA}$ 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 were optimized by hand.
– The approach of Umeyama $\hat{X}_{Ume}$ based on orthogonal relaxation is not as competitive.
– Using the simple 2opt greedy-strategy as post-processing step significantly improves the solution in most cases.

In summary, these results indicate that the convex programming approach $\hat{X}_{QPB}$ embedded in a more sophisticated search strategy (compared to 2opt) is an attractive candidate for solving the weighted graph-matching problem.

Random ground-truth experiments.
We created many problem instances of (1) by randomly computing graphs. The probability that an edge is present in the underlying complete graph was about 0.3. For each pair of graphs, the global optimum was computed using an exact search algorithm.

Table 2 summarizes the statistics of our results. The notation explained in the previous Section was used. The first column on the left shows the problem size n together with the number of random experiments in angular brackets. The number pairs in round brackets denote the number of experiments for which the global optimum was found with/without the 2opt greedy-strategy as a post-processing step. Furthermore the worst case, the best case, and the average case

QAP	X^*	QPB	$\hat{X}_{QPB}$	$\hat{X}_{QPB}+$	$\hat{X}_{GA}$	$\hat{X}_{GA}+$	$\hat{X}_{Ume}$	$\hat{X}_{Ume}+$
chr12c	11156	-22648	20306	15860	19014	11186	40370	11798
chr15a	9896	-48539	26132	14454	30370	11062	60986	17390
chr15c	9504	-47409	29862	17342	23686	13342	76318	13338
chr20b	2298	-7728	6674	2858	6290	2650	10022	3294
chr22b	6194	-20995	9942	6848	9658	6732	13118	7418
esc16b	292	250	296	292	298	292	306	292
rou12	235528	205461	278834	246712	273438	246282	295752	251848
rou15	354210	303487	381016	371480	457908	359748	480352	384018
rou20	725522	607362	804676	746636	840120	738618	905246	765872
tai10a	135028	116260	165364	143260	168096	135828	189852	147838
tai15a	388214	330205	455778	399732	451164	400328	483596	405442
tai17a	491812	415578	550852	513170	589814	505856	620964	526814
tai20a	703482	584942	799790	740696	871480	724188	915144	775456
tai30a	1818146	1517829	1996442	1883810	2077958	1886790	2213846	1875680
tai35a	2422002	1958998	2720986	2527684	2803456	2496524	2925390	2544536
tai40a	3139370	2506806	3529402	3243018	3668044	3249924	3727478	3282284

Table 1. Results of the QAPLIB benchmark experiments (see text).

for the relative values for each of the three estimates presented in Sections 3 and 4 are shown (note that these values are smaller than 1 because the value of the objective function (1) is negative for this class of experiments). In summary, the conclusions with respect to the QAPLIB-experiments are confirmed.

	$\hat{X}_{QPB}/X^*$			$\hat{X}_{Ume}/X^*$			$\hat{X}_{GA}/X^*$		
	mean	worst case	best case	mean	worst case	best case	mean	worst case	best case
n=9 [128]	(22/53)			(7/29)			(31/55)		
	0.87607	0.43552	1	0.638244	0.0651729	1	.948342	.7756129	1
2opt	0.966155	0.79256	1	0.928304	0.753007	1	.9699138	.843046	1
n=11 [42]	(3/11)			(0/7)			(7/10)		
	0.824023	0.514964	1	0.636159	0.295194	0.998591	.940740	.8338586	1
2opt	0.962258	0.842204	1	0.933206	0.811326	1	.9588626	.8434407	1
n=15 [99]	(0/5)			(0/1)			(4/11)		
	0.741563	0.232741	0.938917	0.131333	0.225983	0.863508	.916225	.105164	1
2opt	0.925801	0.777494	1	0.890131	0.74688	1	.9576297	.8205957	1

Table 2. Statistics of the results of random ground-truth experiments (see text).

6 Conclusion

We have shown that, based on advanced techniques from convex optimization theory, suboptimal solutions to the weighted graph-matching problem can be computed which are competitive with respect to recent deterministic annealing approaches. In contrast to annealing approaches, however, the convex approach exhibits two favorable properties: Firstly, no tuning parameters are needed. Secondly, it computes a lower bound and thus can be used as a subroutine within an exact search strategy like branch-and-bound, for example. As a result, it is an attractive candidate for solving matching problems in the context of view-based object recognition.

Acknowledgment: We are thankful for discussions with Prof. Dr.-Ing. W. Förstner, D. Cremers and J. Keuchel.

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. S. Umeyama. An eigendecomposition approach to weighted graph matching problems. *IEEE Trans. Patt. Anal. Mach. Intell.*, 10(5):695–703, 1988.
4. G. Vosselmann. *Relational matching*, volume 628 of *Lect. Not. Comp. Sci.* Springer, 1992.
5. W.K. Konen, T. Maurer, and C. von der Malsburg. A fast dynamic link matching algorithm for invariant pattern recognition. *Neural Networks*, 7(6/7):1019–1030, 1994.
6. S. Gold and A. Rangarajan. A graduated assignment algorithm for graph matching. *IEEE Trans. Patt. Anal. Mach. Intell.*, 18(4):377–388, 1996.
7. A.D.J. Cross, R.C. Wilson, and E.R. Hancock. Inexact graph matching using genetic search. *Pattern Recog.*, 30(6):953–970, 1997.
8. 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.
9. 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.
10. S. Ishii and M. Sato. Doubly constrained network for combinatorial optimization. *Neurocomputing*, 2001. to appear.
11. G. Finke, R.E. Burkard, and F. Rendl. Quadratic assignment problems. *Annals of Discrete Mathematics*, 31:61–82, 1987.
12. 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.
13. R.E. Burkard, S. Karisch, and F. Rendl. Qaplib – a quadratic assignment problem library. *J. Global Optimization*, 10:391–403, 1997.
14. 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.
15. A. Rangarajan, A. Yuille, and E. Mjolsness. Convergence properties of the softassign quadratic assignment algorithm. *Neural Computation*, 11(6):1455–1474, 1999.
16. C. Schellewald, S. Roth, and C. Schnörr. Evaluation of spectral and convex relaxations to the quadratic assignment of relational object views. Comp. science series, technical report, Dept. Math. and Comp. Science, University of Mannheim, Germany, 2001. in preparation.
17. 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.