Solving Quadratically Constrained Geometrical Problems using Lagrangian Duality

Carl Olsson
calle@maths.lth.se
Centre for Mathematical Sciences,
Lund University, Sweden.

Anders Eriksson
anders.eriksson@adelaide.edu.au
School of Computer Science
University of Adelaide, Australia.

Abstract

In this paper we consider the problem of solving different pose and registration problems under rotational constraints. Traditionally, methods such as the iterative closest point algorithm have been used to solve these problems. They may however get stuck in local minima due to the non-convexity of the problem. In recent years methods for finding the global optimum, based on Branch and Bound and convex under-estimators, have been developed. These methods are provably optimal, however since they are based on global optimization methods they are in general more time consuming than local methods.

In this paper we adopt a dual approach. Rather than trying to find the globally optimal solution we investigate the quality of the solutions obtained using Lagrange duality. Our approach allows us to formulate a single convex semidefinite program that approximates the original problem well. [1]

1 Geometric Registration

A frequently occurring and by now a classical problem in computer vision, robotic manipulation and photogrammetry is the registration problem. That is, finding the transformation between two coordinate systems, see [5] and the references therein.

There are a number of solutions proposed and perhaps the most well-known is by Horn *et al.* [6]. They derive a closed-form solution for the Eucledian (or similarity) transformation that minimizes the sum of squares error between the transformed points and the measured points.

The more general problem of finding the registration between two 3-D shapes was considered in [2], where the *iterative closest point* (ICP) algorithm was proposed to solve the problem. The algorithm is able to cope with different geometric primitives, like point sets, line segments and different kinds of surface representations. However, the algorithm requires a good initial transformation in order to converge to the globally optimal solution, otherwise only a local optimum is obtained.

In [8], the method of Horn *et al.* [6] was generalized by incorporating point, line and plane features in a common framework. Given *point-to-point*, *point-to-line*, or *point-to-plane* correspondences, it was demonstrated how the transformation (Euclidean or similarity) relating the two coordinate systems can be computed based on a geometrically meaningful cost-function. The algorithm was based on relaxing the non-convex problem by convex under-estimators and then using branch and bound to focus in on the global solution [1].

In this paper we adopt a dual approach. Rather than trying to find the globally optimal solution we investigate the quality of the solutions obtained using Lagrange duality. Our approach allows us to formulate a single convex semidefinite program that approximates the original problem well. We show on both synthetic and real data that the approximation is very close to the the global optimum, in particular if the noise level is low. Further more, we show that it is possible to obtain lower bounds on the global optimum from our solution using standard duality theory.

Semidefinite programming has previously been used for solving large scale binary quadratic optimization problems [7]. These problems are often NP-complete, however due to a result in [4], semidefinite programming have been shown to produce good approximations while still being a polynomial algorithm. On the other hand for very large problems, polynomial execu-

[1]This work has been supported by the European Commission's Sixth Framework Programme under grant no. 011838 as part of the Integrated Project SMErobotTM and the Swedish Research Council through grant no. 2007-6476 'Efficient Optimization Methods for Large-Scale Problems in Computer Vision'.

tion time might still be too much. In the vision community it has mostly been used for large scale combinatorial problems resulting in long execution times. However, since the matrices that arise in the applications that we consider in this paper are quite small (typically less than 100×100), it provides an efficient way to produce lower estimates and approximative solutions.

2 Registration Problems

In its most general form, the problems we are interested in can be described as geometric registration problems. The objective is to relate measurements in one coordinate frame to an object model in another. In this section we consider the case where the measurements are 3D-points and the model consists of planes in 3D-space. The goal is to find a Eucledian transformation $(R, t) \in SO(3) \times \mathbb{R}^3$ that places the measurement point x_i as close to the corresponding model plane π_i as possible. It was shown in [8], that given a number of planes π_i in one coordinate system and points x_i in another, $i = 1, ..., m$ this problem can be formulated as the optimization problem

$$\mu^* = \min_{\substack{R \in SO(3) \\ t \in \mathbb{R}^3}} \quad \sum_{i=1}^m d(Rx_i + t, \pi_i), \tag{1}$$

where $d(x, \pi)$ denotes the squared distance between point x and the plane π. If we let y_i be an arbitrary point on π_i and n_i be the unit normal, then we can rewrite the problem as

$$\min_{\substack{R \in SO(3) \\ t \in \mathbb{R}^3}} \quad \sum_{i=1}^m ||n_i^T(Rx_i + t - y_i)||_2^2. \tag{2}$$

If we disregard the constraint $R \in SO(3)$ this is a linear least squares problem in the unknowns R and t. However since $R \in SO(3)$ we also have $R^T R = I$ giving the following quadratically constrained problem

$$\min \quad \sum_{i=1}^m ||n_i^T(Rx_i + t - y_i)||_2^2 \tag{3}$$

$$\text{s.t.} \quad R^T R - I = 0 \tag{4}$$

Since $R^T R - I$ is symmetric, equation (4) consists of 6 quadratic constraints.

Next we rewrite the problem in vector form. We let $v^T = [\, r_{11} \; r_{12} \; \dots \; r_{33} \; t_1 \; t_2 \; t_3 \,]$. The objective function is quadratic and (3) can therefore be written

$$v^T A v + 2b^T v, \tag{5}$$

where A and b are determined from the model and measurement data. (To simplify notation we have dropped the constant since this does not change the optimizer.) Note that $A \succeq 0$. By introducing the 12×3-block matrices $E_1 = [\, I \; 0 \; 0 \; 0 \,]$, $E_2 = [\, 0 \; I \; 0 \; 0 \,]$ and $E_3 = [\, 0 \; 0 \; I \; 0 \,]$ we may rewrite the problem as

$$\mu^* = \min \quad v^T A v + 2b^T v \tag{6}$$

$$\text{s.t.} \quad v^T E_i^T E_j v = \delta_{ij} \tag{7}$$

where $\delta_{ij} = 1$ if $i = j$ and 0 otherwise.

This problem is non-convex and therefore difficult to solve without employing complex global optimization techniques (such as in [8, 1]). In this paper we are interested in finding good approximate solutions using methods from convex optimization, which are inherently easier and faster than methods from global optimization. Typically this is done by somehow relaxing the original problem into a convex problem where the solution provides a lower bound on the optimal value μ^* of the original problem. The goal is to find a relaxation that gives a lower bound which is as close to μ^* as possible, or equivalently, a lower bound that is as large as possible. If the lower bound attains μ^* then the relaxation is said to be tight. In this case the exact solution can often be obtained.

2.1 The Linear Relaxation

As stated previously equation (7) consists of 6 quadratic constraints. It is well known that quadratic problems with no more than 2 quadratic constraints can be solved in polynomial time. Hence we are forced to search for approximative solutions. The first relaxation is simply to drop the constraints (7). That is, we solve the problem

$$\mu_l = \min_v v^T A v + 2b^T v. \tag{8}$$

Note that solving (8) corresponds to finding the best affine transformation instead of Euclidian transformation. A typical way of generating starting points for local methods is to find the best affine transformation and then to upgrade it to a similarity transformation using singular value decomposition. This technique works well if the noise level is low and the number of measurements is high.

Since the set of feasible matrices is larger if the rotation constraint is disregarded we can conclude that the minimum μ_l fulfills

$$\mu_l \le \mu^*. \tag{9}$$

(The matrix that minimizes (6) with (7) is also feasible in (8).) Hence the least squares solution provides a lower bound on the minimum value of the original problem.

A nice property of the linear relaxation is that for a large number of measurements it will be roughly correct. If the measurements are corrupted by Gaussian noise then (8) will be the statistically optimal estimation for the problem with affine transformations. Since the set of Euclidian transformations is a subset of the affine transformations it is easy to see that as the number of measurements grow the affine relaxation will approach a correct solution despite having dropped the rotation constraints.

2.2 The Lagrangian Relaxation

As discussed in the previous section the problem is easily solved if the constraints are dropped. This gives us a lower bound on the objective value. However we can easily obtain other lower bounds by adding $\mu_{ij}(v^T E_i^T E_j v - \delta_{ij})$ to the objective function (since in the original problem these terms should be zero). For each different μ_{ij} we may obtain a linear relaxation. Lagrangian duality can be viewed as finding the relaxation that gives the largest objective value. The Lagrangian of (6) becomes

$$L(v,\lambda) = v^T A v + 2b^T v + \sum_{i,j} \lambda_{ij}(v^T E_{ij} v - \delta_{ij}) = v^T \left(A + \sum_i \lambda_{ij} E_{ij}\right) v + 2b^T v - \sum_{i,j} \lambda_{ij}\delta_{ij}, \quad (10)$$

where $E_{ij} = \frac{1}{2}(E_i^T E_j + E_j^T E_i)$. Thus finding the largest relaxation can be written as

$$\mu_L = \max_\lambda \min_v L(v,\lambda) \quad (11)$$

Since the linear relaxation corresponds to $\lambda = 0$ it is easy to see that

$$\mu_l \leq \mu_L \leq \mu^*. \quad (12)$$

That is, the Lagrangian relaxation is always at least as good as the linear relaxation. In practice it is usually much better (see section 3).

For a fixed λ the inner minimization of (11) is only finite if the matrix $A + \sum_{i,j} \lambda_{ij} E_{ij}$ is positive semidefinite. In this case the minimizer can be computed analytically by taking the gradient with respect to v, which yields

$$v = -(A + \sum_{i,j} \lambda_{ij} E_{ij})^{-1} b \quad (13)$$

Inserting into (10) we see that (11) can be written

$$\begin{aligned} \max_\lambda \quad & -b^T \left(A + \textstyle\sum_{i,j} \lambda_{ij} E_{ij}\right)^{-1} b + \textstyle\sum_{ij} \delta_{ij}\lambda_{ij} \\ \text{s.t} \quad & A + \textstyle\sum_{i,j} \lambda_{ij} E_{ij} \succeq 0 \end{aligned} \quad (14)$$

The above problem is the dual program when the primal variables have been eliminated. It can be shown that the objective function is concave and the constraint is convex, hence this program can be solved efficiently. However to be able to use a standard solver, such as SeDuMi [9] we need to write it as a standard linear semidefinite program. This can be done using the Schur-complement (see [3]). First we add the artificial variable γ and rewrite the program as

$$\max_{\gamma,\lambda} \quad \gamma \quad (15)$$

$$\text{s.t.} \quad A + \textstyle\sum_{i,j} \lambda_{ij} E_{ij} \succeq 0 \quad (16)$$

$$-b^T (A + \sum \lambda_{ij} E_{ij})^{-1} b + \sum \delta_{ij}\lambda_{ij} - \gamma \geq 0 \quad (17)$$

Let

$$A(\lambda,\gamma) = \begin{bmatrix} \sum_{ij} \delta_{ij}\lambda_{ij} - \gamma & b^T \\ b & A + \sum_{i,j} \lambda_{ij} E_{ij} \end{bmatrix} \quad (18)$$

Using the Schur complement [3], the dual problem can finally be written as a linear semidefinite program

$$\max_{\gamma,\lambda} \quad \gamma \quad (19)$$

$$\text{s.t.} \quad A(\lambda,\gamma) \succeq 0. \quad (20)$$

This program is usually referred to the dual semidefinite program. Using the same technique as above it is possible to derive the primal semidefinite program

$$\min_{V \succeq 0} \quad \text{tr}\left(\begin{bmatrix} A & b \\ b^T & 0 \end{bmatrix} V\right) \quad (21)$$

$$\text{s.t.} \quad \text{tr}(E_{ij} V) = \delta_{ij} \quad (22)$$

$$\text{tr}(\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} V) = 1. \quad (23)$$

Note that if V^* solves the primal problem and V^* can be written

$$V^* = \begin{bmatrix} v^*(v^*)^T & v^* \\ (v^*)^T & 1 \end{bmatrix} \quad (24)$$

then v^* is the global optimum of problem (6).

When using a primal-dual interior point as SeDuMi one obtains both a solution to the primal and the dual programs. To generate good approximate solutions we will take the eigenvector corresponding to the largest eigenvalue of the primal solution V as our approximate solution. If V is close to a rank one matrix this should be a good approximation.

2.3 Generalization to other types of Correspondences.

In the previous section we derived our method in the case of point-to-plane correspondences only, however it is easily generalized to point-to-point and point-to-plane correspondences.

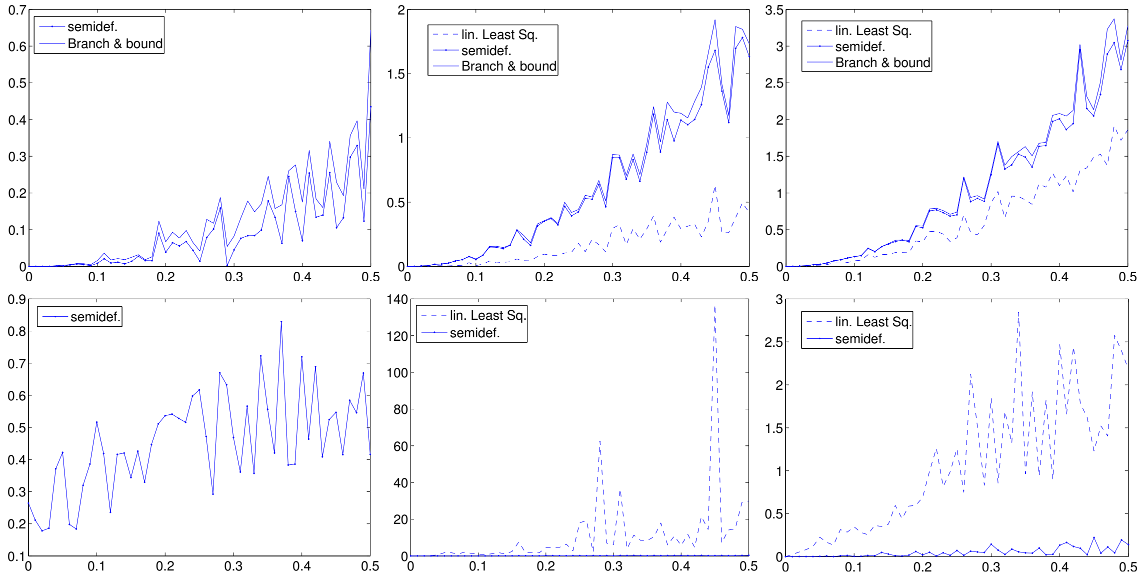

Figure 1. Top: The objective value of the obtained solution as a function of the noise level for three different problem sizes, m = 7, 14 and 20. For comparison we also plot the global minimum, given by the Branch & Bound method, and the Linear Least Squares solution. Bottom: The value of $||R^TR - I||_2^2$ for the obtained solution as a function of the noise level.

Let x_i^l be the measured 3D points and let l_i, $i = 1, ..., m_l$, be the corresponding lines. Then the sum of squared distances between the transformed points and the lines can be written

$$\sum_{i=1}^{m_l} ||(I - v_i v_i^T)(Rx_i^l + t - y_i^l)||_2^2, \qquad (25)$$

where v_i is a unit direction vector for the line l_i and y_i^l is any point on the line l_i. Since this is a quadratic function it can also be written as (5).

The case of point-to-point correspondences is the easiest one. Let x_i^p be the measured points and y_i^p be the corresponding points $i = 1, ..., m_p$. The objective function can in the same way as for the point-to-line case be written as

$$\sum_{k=1}^{3} \sum_{i=1}^{m_p} ||e_k(Rx_i^p + t - y_i^p)||^2 \qquad (26)$$

where e_k is the k'th row of the identity matrix. Again this can be written in the form (5).

3 Experiments

In this section we present a few experiments that shows that even though the program we have derived is a relaxation it provides surprisingly good solutions.

In our first experiment we use synthetic data to test the performance for various noise levels. The data was created by randomly placing m planes such that they intersect the sphere centered around zero with radius 10. Then we randomly selected one point from each plane and added noise of standard deviation σ. For each noise level we generated 10 experiments and plotted the mean result of the various methods.

The first row of figure 1 shows the objective value of the obtained solution for different noise levels and different problem sizes. The relaxation can be seen to approximate the global optimum well, particularly for low noise level. Also it is a considerable improvement compared to the linear least squares solution. Note that for the case of 7 point-to-plane correspondences there is not enough data to compute a least squares solution. Also note that as the number of measurements increases the least squares solution will approach the correct solution. This is to be expected since the measurements are corrupted with normally distributed noise. The second row of figure 1 shows the value of $||R^TR - I||_2^2$ for the approximate solutions. The semidefinite solution appears to be quite close to being a rotation matrix while the linear least squares is far off, although as expected its performance improves as the number of measurements increases. Note that the y-scale is not the same in all pictures.

For our next experiment we used real data. Figure 2 shows the setup for this experiment. We used

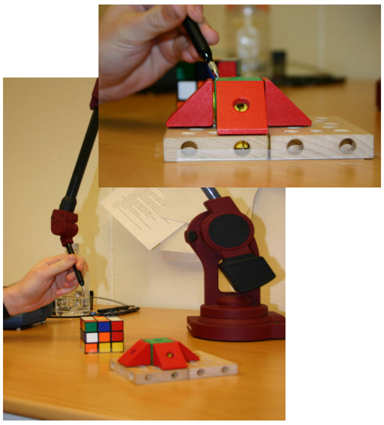

Figure 2. The experimental setup for the spacestation experiment.

a MicroScribe-3DLX 3d scanner to measure the 3D-coordinates of some points on the toy shown in figure 2. (By request of the designer we will refer to the toy model as the space station.) The 3D-scanner consists of a pointing arm with five degrees of freedom which measures 3D-coordinates. In total we measured 49 points on the toy model visible in figure 2. We created a computer model consisting of planes, lines and points (see figure 3).

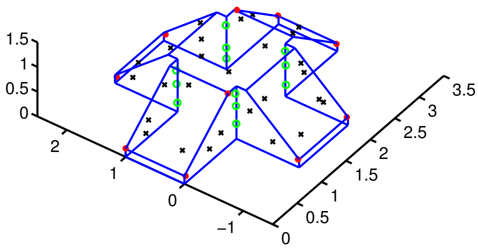

Figure 3. The model of the spacestation and the optimal registration.

Among the measured points 27 where known to belong to a certain plane in the model, 12 to a line and 10 to point. Note that we only considered point-to-plane correspondences in the previous section, however it is easy to extend to other correspondences as well. Figure 3 shows the model and the resulting registration. The points marked with black crosses are points measured on the planes, the points marked with green rings are measured on lines and the points marked with red stars are measured on corners. In table 1 we show the total error for the different correspondences and the running times. For comparison we have also included the errors and running times obtained with the other methods. In this case the solution of the primal semidefinite program has rank one and is therefore the optimal solution (same as Branch & Bound). Note that Horn's method only measures the point-to-point errors resulting in a higher total error.

Residuals:	*B & B*	*Horn*	*Least Sq.*	*Semidef*
point-point	0.0083	0.0063	0.0221	0.0083
point-line	0.0018	0.0036	0.0015	0.0018
point-plane	0.0046	0.0098	0.0046	0.0046
Total	0.0147	0.0197	0.0282	0.0147
Run-time:	24	0.062	0.0056	0.46

Table 1. Resulting reconstruction errors for the space station problem and execution time in seconds when using the different methods.

References

[1] S. Agarwal, M. K. Chandraker, F. Kahl, D. J. Kriegman, and S. Belongie. Practical global optimization for multiview geometry. In *Proc. 9th European Conf. on Computer Vision, Graz, Austria*, pages 592–605, 2006.

[2] P. Besl and N. McKay. A method for registration two 3-d shapes. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 14(2):232–256, 1992.

[3] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[4] M. Goemans and D. Wiliamson. Improved approximation algorithms for maximum cut and satisfiability problem using semidefinite programming. *J.ACM*, 42(6):1115–1145, 1995.

[5] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.

[6] B. K. Horn, H. M. Hilden, and S. Negahdaripour. Closed-form solution of absolute orientation using ortonormal matrices. *Journal of the Optical Society of America A*, 5, 1988.

[7] J. Keuchel, C. Schnörr, C. Schellewald, and D. Cremers. Binary partitioning, perceptual grouping, and restoration with semidefinite programming. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(11):1364–1379, 2006.

[8] C. Olsson, F. Kahl, and M. Oskarsson. The registration problem revisited: Optimal solutions from points lines and planes. In *International Conference on Computer Vision and Pattern Recognition*, New York City, USA, 2006.

[9] J. F. Sturm. Using sedumi 1.02, a matlab toolbox for optimization over symmetric cones. 1998.