

Autocalibration via Rank-Constrained Estimation of the Absolute Quadric

Manmohan Chandraker

University of California, San Diego

Sameer Agarwal

University of Washington, Seattle

Fredrik Kahl

Lund University, Sweden

David Nistér

Microsoft Research, Seattle

David Kriegman

University of California, San Diego

Abstract

We present an autocalibration algorithm for upgrading a projective reconstruction to a metric reconstruction by estimating the absolute dual quadric. The algorithm enforces the rank degeneracy and the positive semidefiniteness of the dual quadric as part of the estimation procedure, rather than as a post-processing step. Furthermore, the method allows the user, if he or she so desires, to enforce conditions on the plane at infinity so that the reconstruction satisfies the chirality constraints.

The algorithm works by constructing low degree polynomial optimization problems, which are solved to their global optimum using a series of convex linear matrix inequality relaxations. The algorithm is fast, stable, robust and has time complexity independent of the number of views. We show extensive results on synthetic as well as real datasets to validate our algorithm.

1. Introduction

By its very nature, the problem of determining the internal parameters of a camera from image data alone has generated tremendous theoretical interest in the field of computer vision since its introduction in [4]. The appeal of the problem is practical as well, since a generic self-calibration method eliminates the need for a special offline procedure, such as using a calibration grid. While a projective reconstruction of the scene can be computed from image coordinates alone, the goal of autocalibration is to compute the projective transformation (homography) that upgrades the projective reconstruction to a metric reconstruction.

The basic tenet of autocalibration is the constancy of the absolute conic under rigid body motion of the camera. This is encoded conveniently in the absolute dual quadric (\mathbf{Q}_∞^*) formulation of autocalibration [23]. Once estimated, an eigenvalue decomposition of the dual quadric yields the homography that relates the projective reconstruction to a Euclidean reconstruction.

The common practice in estimating \mathbf{Q}_∞^* is to enforce its rank degeneracy as a post-processing step by simply dropping the smallest singular value. The rank degeneracy of the absolute quadric has an important physical interpretation: enforcing it is equivalent to demanding a common support plane for the absolute conic over the multiple views. That, indeed, is the real advantage that an absolute quadric based method affords over one based on, say, the Kruppa constraints. Thus, for more than two views, it does not make sense to estimate \mathbf{Q}_∞^* without enforcing the rank condition.

In this paper, we propose a method for estimating the absolute quadric where its rank deficiency is imposed within the estimation procedure. A significant drawback of several prior approaches for autocalibration is the difficulty in ensuring a positive (or negative, depending on scale) semidefinite \mathbf{Q}_∞^* . As has been discussed in the literature [9], it is usually not correct to simply output the closest positive semidefinite matrix as a post-processing step since it might lead to a degenerate calibration. Our formulation explicitly demands a positive (or negative) semidefinite \mathbf{Q}_∞^* as an output of the optimization problem, which ensures that the resulting dual image of the absolute conic (DIAC) can be decomposed into its Cholesky factors to yield the calibration parameters of the cameras.

It is well-established that the principal difficulty in autocalibration lies in the affine upgrade step, which involves a precise estimation of the plane at infinity, π_∞ . In our approach, estimating the plane at infinity is encapsulated in the absolute quadric estimation itself, as π_∞ lies in the null-space of \mathbf{Q}_∞^* . Further, it is reasonable to demand that chirality holds, that is, the points and camera centers lie on one side of the plane at infinity [7]. It has been argued in recent literature [16] that it is most important for any reconstruction method to start by satisfying the requirements of chirality, in particular, with respect to the camera centers. Intuitively, moving across the plane at infinity requires jumps across large “basins” in the search space. Thus, imposing chirality constraints increases the chances of an autocalibration routine to start in the correct region of the search space.

Given reasonable assumptions on the internal parameters of the camera, such as zero skew and unit aspect ratio, we pose the problem of estimating the positive semidefinite, rank-degenerate absolute quadric as one of minimizing a polynomial objective function, subject to polynomial equality and inequality conditions. Chirality conditions translate into polynomial inequality constraints on the entries of the absolute quadric, so they can be readily included in the polynomial system that we seek to minimize. Our formulation allows us to compute the global minimum for such a polynomial system using a series of convex linear matrix inequality relaxations [13, 14].

In summary, the contributions of this paper are:

- We present a fast and reliable method for autocalibration by estimating the absolute dual quadric, where its rank degeneracy and positive semidefiniteness are imposed within our optimization framework.
- We can demand that our reconstruction satisfy the requirements of chirality by imposing constraints on the plane at infinity during our estimation procedure.
- We globally minimize a reasonable objective function based on camera matrices alone to deduce the entries of the absolute quadric subject to all the above constraints.

2. Background

Unless stated otherwise, we will denote 3D points by homogeneous 4-vectors (such as $\mathbf{X} = (X_1, X_2, X_3, X_4)^\top$) and 2D points by homogeneous 3-vectors (such as $\mathbf{x} = (x_1, x_2, x_3)^\top$). A projective camera is represented as $\mathbf{P} = \mathbf{K}[\mathbf{R}|\mathbf{t}]$, where the intrinsic calibration parameters of the camera are encoded in the upper triangular matrix \mathbf{K} and (\mathbf{R}, \mathbf{t}) denote the exterior orientation of the camera. One of the objectives in autocalibration is to recover the matrix \mathbf{K} , which we parametrize as

$$\mathbf{K} = \begin{bmatrix} f_x & s & u \\ 0 & f_y & v \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

where f_x, f_y stand for the focal lengths in the x and y directions, s denotes skew and (u, v) the position of the principal point.

Suppose we have a projective reconstruction $\{\mathbf{P}^i, \mathbf{X}^j\}$ for $i = 1, \dots, m$ and $j = 1, \dots, n$. We wish to determine the projective transformation \mathbf{H} that takes the projective reconstruction back to Euclidean $\{\mathbf{P}_M^i, \mathbf{X}_M^j\}$ where

$$\begin{aligned} \mathbf{P}_M^i &= \mathbf{P}^i \mathbf{H}, \quad i = 1, \dots, m \\ \mathbf{X}_M^j &= \mathbf{H}^{-1} \mathbf{X}^j, \quad j = 1, \dots, n. \end{aligned} \quad (2)$$

We can always perform the projective reconstruction such that $\mathbf{P}^1 = [\mathbf{I}|\mathbf{0}]$. We can choose the world Euclidean frame

to coincide with the first camera, so that $\mathbf{P}_M^1 = \mathbf{K}^1 [\mathbf{I}|\mathbf{0}]$. Then, the form of \mathbf{H} can be deduced as

$$\mathbf{H} = \begin{bmatrix} \mathbf{K}^1 & \mathbf{0} \\ -\mathbf{p}^\top \mathbf{K}^1 & 1 \end{bmatrix} \quad (3)$$

where the plane at infinity in the projective reconstruction is $\pi_\infty = (\mathbf{p}, 1)^\top$. The canonical form of π_∞ is $(0, 0, 0, 1)^\top$ and is moved by a projective transformation.

The absolute conic, Ω_∞ , is the point conic on π_∞ governed by $X_1^2 + X_2^2 + X_3^2 = 0, X_4 = 0$. A principal quantity of interest in autocalibration is the dual image of the absolute conic (DIAC): $\omega^* = \mathbf{K}\mathbf{K}^\top$. Clearly, once the DIAC is estimated, it is only a matter of computing its Cholesky decomposition to recover \mathbf{K} . (Of course, provided that the recovered ω^* is positive semidefinite!) We refer the reader to [9] for more details and turn our attention to a more powerful geometric entity useful for self-calibration.

2.1. The Absolute Dual Quadric

The absolute dual quadric is the dual to the absolute conic Ω_∞ in metric 3-space and has the form

$$\mathbf{Q}_\infty^* = \tilde{\mathbf{I}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0}^\top & 0 \end{bmatrix}. \quad (4)$$

Any plane in the envelope of the absolute dual quadric is tangent to the absolute conic Ω_∞ .

The following are known properties of \mathbf{Q}_∞^* which are used in this paper. We will simply state them here without proof: \mathbf{Q}_∞^* is a degenerate conic - it is singular, of rank 3. \mathbf{Q}_∞^* is symmetric, positive (or negative, depending on sign) semi-definite. Under a transformation \mathbf{H} , \mathbf{Q}_∞^* transforms as $\mathbf{Q}_\infty^{*\prime} = \mathbf{H}\mathbf{Q}_\infty^* \mathbf{H}^\top$. \mathbf{Q}_∞^* is fixed under a similarity transformation, but is moved by a general projective transformation. The plane at infinity, π_∞ , is a null vector of \mathbf{Q}_∞^* .

$$\mathbf{Q}_\infty^* \pi_\infty = \mathbf{0}. \quad (5)$$

2.2. Autocalibration Using the Absolute Quadric

The projective transformation that takes the projective reconstruction back to Euclidean is the one that takes \mathbf{Q}_∞^* measured in the projective reconstruction back to its canonical form $\tilde{\mathbf{I}}$, that is, $\mathbf{Q}_\infty^* = \mathbf{H}\tilde{\mathbf{I}}\mathbf{H}^\top$. For a camera matrix \mathbf{P}^i , the dual image of the absolute conic, ω^{*i} is the image under \mathbf{P}^i of the dual of the absolute conic Ω_∞ , that is the image of \mathbf{Q}_∞^* under \mathbf{P}^i . So,

$$\omega^{*i} = \mathbf{P}^i \mathbf{Q}_\infty^* \mathbf{P}^{i\top}. \quad (6)$$

Now, we can impose constraints on the entries of ω^{*i} to constrain the entries of \mathbf{Q}_∞^* . Once \mathbf{Q}_∞^* is determined, \mathbf{H} can be recovered by a simple eigendecomposition.

2.3. Chirality

Chirality constraints demand that the reconstruction satisfy a very basic criterion: the imaged scene points must be in front of the camera [7]. A general projective transformation need not preserve the convex hull of a point set. That is, the scene can be split across the plane at infinity in a projective reconstruction. A quasi-affine reconstruction is one that differs from the Euclidean scene by a projective transformation, but in which the plane at infinity is guaranteed not to split the point set. A quasi-affine reconstruction can be computed from the solution of the so-called ‘‘chiral inequalities’’ determined by all the scene points and camera centers [7]. For a thorough introduction to the concept, we refer the reader to the corresponding chapter in [9].

2.4. Polynomial Minimization using LMI Relaxations

We are interested in solving polynomial optimization problems of the following class:

$$\begin{aligned} \min \quad & f(\mathbf{x}) \\ \text{subject to} \quad & g_i(\mathbf{x}) \geq 0, \quad i = 1, 2, \dots, m, \end{aligned} \quad (7)$$

where $f(\mathbf{x})$ and $g_i(\mathbf{x})$ are scalar polynomials of the vector indeterminate $\mathbf{x} \in \mathbf{R}^n$. Let p^* denote the minimum objective value (if it exists) of the above problem. Then, a convex relaxation is, by construction, a convex optimization problem with minimum objective value p_r^* such that $p_r^* \leq p^*$. Hence, by solving the relaxed problem, a lower bound on the original objective function is obtained.

When optimizing a scalar polynomial objective function subject to polynomial constraints, convex relaxations can be obtained by gradually adding lifting variables and constraints corresponding to linearizations of monomials up to a given degree. This is the technique we will adopt. The LMI relaxation covering monomials up to a given even degree 2δ is referred to as the LMI relaxation of order δ . The standard Shor relaxation in mathematical programming [21] can be regarded as a first-order LMI relaxation.

For an LMI relaxation of order δ , let $\mathbf{v}_\delta(\mathbf{x})$ be a vector containing all monomials up to degree δ including the constant term 1. To form the relaxed optimization problem of the original problem (7), the following steps are required:

1. Linearize the objective function $f(\mathbf{x})$ by lifting: $x_1^{k_1} x_2^{k_2} \dots x_n^{k_n}$ is replaced with the new lifting variable $y_{k_1 k_2 \dots k_n}$. Thus, the linearized objective function can be written $\mathbf{f}^\top \mathbf{y}$ for a constant coefficient vector \mathbf{f} and the lifting vector \mathbf{y} .
2. Apply lifting to the LMI constraint $g_i(\mathbf{x}) \mathbf{v}_{\delta-1}(\mathbf{x}) \mathbf{v}_{\delta-1}(\mathbf{x})^\top \succeq 0$ for each constraint $g_i(\mathbf{x}) \geq 0$. Denote the linearized constraint by $M_{\delta-1}(g_i(\mathbf{y})) \succeq 0$.

3. Add the LMI moment matrix constraint which corresponds to linearizing the trivial constraint $\mathbf{v}_\delta(\mathbf{x}) \mathbf{v}_\delta(\mathbf{x})^\top \succeq 0$. Denote the linearized constraint by $M_\delta(\mathbf{y}) \succeq 0$.

To summarize, the following SDP is solved for the LMI relaxation of order δ of the problem (7):

$$\begin{aligned} \min \quad & \mathbf{f}^\top \mathbf{y} \\ \text{subject to} \quad & M_{\delta-1}(g_i(\mathbf{y})) \succeq 0, \quad i = 1, 2, \dots, m, \\ & M_\delta(\mathbf{y}) \succeq 0. \end{aligned} \quad (8)$$

If the feasible set $\{\mathbf{x} \mid g_i(\mathbf{x}) \geq 0, i = 1, \dots, m\}$ is compact and under some mild additional assumptions akin to qualification constraints in mathematical programming, it is shown in [14] that the hierarchy of relaxations converges asymptotically $\lim_{\delta \rightarrow \infty} p_\delta^* = p^*$. It turns out that for many of the non-convex polynomial optimization problems, global optima are reached at a given accuracy for a moderate number of lifting variables and constraints, hence for an LMI relaxation of moderate order. A *sufficient condition* for reaching the global optimum is that the moment matrix $M_\delta(\mathbf{y}^*)$ has rank one at the optimum \mathbf{y}^* .

If the solution to the relaxed problem is not tight, that is, $p_\delta^* < p^*$, then an approximate solution may be obtained by simply keeping the lifting variables corresponding to first-order moments. A Matlab toolbox for LMI relaxations can be found in [10].

3. Related Work

There is a significant body of literature within computer vision that deals with autocalibration, beginning with the introduction of the concept in [4]. Approaches to autocalibration can be broadly classified as stratified and direct. The former is a two-step process, whereby the first step involves estimating the plane at infinity for an upgrade to the affine stratum and the metric upgrade is typically performed by estimating \mathbf{K} in a subsequent step.

Estimating the plane at infinity to achieve an affine upgrade is considered the most difficult step in autocalibration [8]. The plane at infinity itself has proven to be a rather elusive entity to estimate precisely. A prior approach has been to exhaustively compute all 64 solutions to the modulus constraints [19], although only 21 of them are physically realizable [20]. An alternate approach involves solving a linear program arising from chirality constraints imposed on the points and camera centers to delineate the region in \mathbf{R}^3 where the first three coordinates of the plane at infinity parametrized as $\pi_\infty = (\mathbf{p}, 1)^\top$ must lie. Subsequently, \mathbf{p} is recovered by a brute force search within this region [8]. In our approach, estimating the plane at infinity is encapsulated in the absolute quadric estimation itself, as π_∞ lies in the null-space of \mathbf{Q}_∞^* .

A variety of linear methods exists for estimating \mathbf{K} for the metric upgrade step, see [9] for more discussion. A drawback of linear approaches is that they do not enforce positive semidefiniteness of the DIAC. One work that the authors are aware of, where estimation of the DIAC is constrained to be positive semidefinite, is [2].

The class of direct approaches to autocalibration are those that directly compute the metric reconstruction from a projective one by estimating the absolute conic. Kruppa equations are view-pairwise constraints on the projection of the absolute quadric. Methods based on the Kruppa equations (or the fundamental matrix), such as [15], are known to suffer from additional ambiguities when used for autocalibration with three or more views [22].

The absolute quadric was introduced as a device for autocalibration in [11, 23], as it is a convenient representation for both the absolute conic and the plane at infinity. Moreover, it can simultaneously be estimated over multiple views. Constraints on the DIAC can be transferred to those on \mathbf{Q}_∞^* using the (known) cameras in the projective reconstruction. The actual solution methods proposed in [23] are a linear approach and one based on sequential quadratic programming. Linear initializations for estimating the dual quadric are also discussed in [17]. It is known that these methods which do not ensure positive semidefiniteness are liable to perform quite poorly with noisy data.

While it can be shown that zero skew alone is sufficient for a metric reconstruction [12], a practitioner must use as much of the available information as possible [9]. We will adopt this latter philosophy and make the safe assumptions that the skew is close to zero, the principal point close to origin and the aspect ratio is close to unity. We allow the focal length to vary to account for varying zoom.

Very recently, there has been interest in developing globally optimal solutions to several problems in multiview geometry. A number of simpler problems in multiview geometry can be formulated in terms of systems of polynomial inequalities [13], which can be globally minimized using the theory of convex linear matrix inequality relaxations [14]. The inherent fractional nature of multiview geometry problems is exploited in [1] to compute the globally optimal solution to triangulation and resectioning, with a certificate of optimality. A branch-and-bound method is used for autocalibration in [5], however, their problem is formulated in terms of the fundamental matrix of view pairs and does not scale beyond a small number of views.

4. Problem Formulation

Many self-calibration algorithms, such as [15, 5], do not estimate the absolute quadric directly. As compared to algorithms such as [23, 16] which estimate the absolute quadric, our main contribution is that we constrain \mathbf{Q}_∞^* to be positive semidefinite and rank degenerate within our optimization

framework. Our optimization problems themselves have been designed with the intention of extracting a globally minimal solution.

Further, we give the user the option to impose the requirements of chirality within the same unified framework, if needed. It has been argued in prior literature that it is desirable to ensure that chirality is satisfied only with respect to the set of camera centers [16]. The reason is that cameras are estimated using robust techniques from several points, so they exhibit better statistical properties than the points themselves. And a few outliers in the scene points can make it impossible to satisfy the requirements of full chirality. Similar to the quasi-affine reconstruction with respect to camera centers (QUARC) in [16], a pairwise twist test ensures that the plane at infinity cannot violate the line segment joining any pair of camera centers. By subsequently imposing the condition on our reconstruction that the plane at infinity lie on one side of all the camera centers, we are guaranteed to recover a \mathbf{Q}_∞^* consistent with the requirements of chirality.

4.1. Imposing rank degeneracy and positive semidefiniteness of \mathbf{Q}_∞^*

Let us suppose an appropriate objective function, $f(\mathbf{Q}_\infty^*)$, has been defined, which depends on the parameters of \mathbf{Q}_∞^* and imposes some desired property on the metric reconstruction. In addition, we demand that the absolute quadric be rank deficient and positive semidefinite. Thus, our optimization problem is of the form:

$$\begin{aligned} \min \quad & f(\mathbf{Q}_\infty^*) \\ \text{subject to} \quad & \text{rank}(\mathbf{Q}_\infty^*) < 4, \quad \mathbf{Q}_\infty^* \succeq 0. \end{aligned} \quad (9)$$

Since \mathbf{Q}_∞^* is a symmetric matrix, it can be parameterized using 10 variables. The condition of rank degeneracy can be imposed by demanding that $\det \mathbf{Q}_\infty^* = 0$, which is a polynomial of degree 4. The positive semidefiniteness of \mathbf{Q}_∞^* can be ensured by asserting that each principal minor of \mathbf{Q}_∞^* have a non-negative determinant, which are polynomial equations of degree at most 3. Thus, an equivalent problem is:

$$\begin{aligned} \min \quad & f(\mathbf{Q}_\infty^*) \\ \text{subject to} \quad & \det(\mathbf{Q}_\infty^*) = 0 \\ & \det[\mathbf{Q}_\infty^*]_{jk} \geq 0, \quad j = 1, 2, 3 \quad k = 1, \dots, \binom{4}{j} \\ & \|\mathbf{Q}_\infty^*\|_F^2 = 1. \end{aligned} \quad (10)$$

where $[\mathbf{Q}_\infty^*]_{jk}$ stands for the k -th $j \times j$ principal minor of \mathbf{Q}_∞^* . Note that one need not impose all of the above inequalities for ensuring semidefiniteness, but doing so may strengthen the convex LMI relaxation. Since \mathbf{Q}_∞^* is only defined up to a scale factor, we use the last equality constraint to fix its norm to one. The above is a system of polynomial

(in)equalities and our objective function is also a polynomial, which can be minimized globally using the theory in Section 2.4.

4.2. Imposing chirality constraints

Next, to impose chirality constraints, recall that the plane at infinity is the null-vector of \mathbf{Q}_∞^* and can be expressed (up to scale) as

$$\pi_\infty = \left[\det(\widehat{\mathbf{Q}}_{\infty 1}^*), \det(\widehat{\mathbf{Q}}_{\infty 2}^*), \det(\widehat{\mathbf{Q}}_{\infty 3}^*), \det(\widehat{\mathbf{Q}}_{\infty i}^*) \right]^\top$$

where $\widehat{\mathbf{Q}}_{\infty i}^*$ represents the 3×3 matrix formed by eliminating the fourth row and the i -th column of \mathbf{Q}_∞^* . The camera center is determined as

$$\mathbf{C}^i = \left[\det(\widehat{\mathbf{P}}^i_1), \det(\widehat{\mathbf{P}}^i_2), \det(\widehat{\mathbf{P}}^i_3), \det(\widehat{\mathbf{P}}^i_4) \right]$$

where $\widehat{\mathbf{P}}^i_j$ stands for the i -th camera in the projective reconstruction with its j -th column eliminated.

Now, chirality constraints are of the form $\pi_\infty^\top \mathbf{C}^i > 0$, which is simply a polynomial of degree 3. Thus, even chirality constraints can be included as part of the polynomial system in our formulation for autocalibration.

4.3. Choice of objective function

Finally, we address the question of the objective function. Over the years, several different objective functions have been specified for autocalibration. The trade-off in designing a suitable objective function, as discussed in [16], is between retaining geometric meaningfulness and ensuring optimality of the recovered solution. We have looked at various choices of the objective function within our polynomial optimization framework, which are described below.

In most situations, it is quite correct to assume that the skew is close to zero and aspect ratio close to unity, while a simple transformation of the image coordinates sets the principal point to $(0, 0)$. Let us assume we have enough prior knowledge of the camera intrinsic parameters in our motion sequence to apply a suitable transformation to the coordinate system that brings the intrinsic parameter matrices of the cameras close to identity. With this ‘‘pre-conditioning’’ based on prior knowledge, we can demand that an algebraic condition be satisfied by the entries of the DIAC so that \mathbf{K} has a form $\text{diag}(f, f, 1)$. One such objective function to be minimized is:

$$f(\mathbf{Q}_\infty^*) := \sum_i (\omega_{11}^{*i} - \omega_{22}^{*i})^2 + \omega_{12}^{*i 2} + \omega_{13}^{*i 2} + \omega_{23}^{*i 2}. \quad (11)$$

Recall that $\omega^{*i} = \mathbf{P}^i \mathbf{Q}_\infty^* \mathbf{P}^{i\top}$, thus $\omega_{jk}^{*i} = \mathbf{p}_j^{\top} \mathbf{Q}_\infty^* \mathbf{p}_k^i$ where, \mathbf{p}_k^i stands for the k -th row of the i -th camera matrix.

Experiments with synthetic data for this objective function are described in Section 5 and results are tabulated in

Table 1. This works well even when we allow focal length to vary such that \mathbf{K} is significantly different from identity. Experimental results for this scenario are also given in Table 1.

A problem with the above objective function is that it is not normalized to account for the scale invariance of ω^* . This can be achieved by dividing each quantity by, say ω_{33}^{*i} , to obtain the following objective function:

$$f(\mathbf{Q}_\infty^*) := \sum_i \frac{(\omega_{11}^{*i} - \omega_{22}^{*i})^2 + \omega_{12}^{*i 2} + \omega_{13}^{*i 2} + \omega_{23}^{*i 2}}{\omega_{33}^{*i 2}}. \quad (12)$$

This is a rational objective function, which can be tackled in our polynomial optimization set-up by introducing a new variable corresponding to each view. The optimization problem we address now has the form:

$$\begin{aligned} \min \quad & \sum_{i=1}^n t_i \\ \text{subject to} \quad & \omega_{33}^{*i} t_i = (\omega_{11}^{*i} - \omega_{22}^{*i})^2 + \omega_{12}^{*i 2} \\ & \quad \quad \quad + \omega_{13}^{*i 2} + \omega_{23}^{*i 2} \\ & \det(\mathbf{Q}_\infty^*) = 0, \quad \mathbf{Q}_\infty^* \succeq 0. \end{aligned} \quad (13)$$

The number of variables in the above optimization problem increases linearly with the number of views. So, it can be solved only for a relatively smaller number of views (three, maybe four) with the current state of the art in polynomial minimization.

It is easy to impose further constraints on the problem, such as zero skew, which corresponds to a quadratic polynomial:

$$\omega_{12}^{*i} \omega_{33}^{*i} = \omega_{13}^{*i} \omega_{23}^{*i}. \quad (14)$$

Other constraints such as principal point at the origin and unit aspect ratio can be similarly imposed as linear or quadratic polynomial constraints.

There can be other principled approaches to formulating a polynomial objective function for autocalibration, assuming constant intrinsic parameters, such as

$$f(\mathbf{Q}_\infty^*) := \sum_i \|\omega^* - \lambda_i \mathbf{P}^i \mathbf{Q}_\infty^* \mathbf{P}^{i\top}\|. \quad (15)$$

For a projective reconstruction such that $\mathbf{P}^1 = [\mathbf{I} \mid \mathbf{0}]$, \mathbf{Q}_∞^* can be parametrized in terms of ω^* and the three parameters of the plane at infinity, which leads to $10 + n$ variables for the n -view problem. Our experimental sections will discuss only the results obtained by using objective functions in (11) and (13).

5. Experiments with synthetic data

We have subjected our algorithm to extensive simulations with synthetic, noisy data to give the reader an idea of its

performance in statistical terms. The implementation is done in Matlab with the GloptiPoly toolbox [10]. An LMI relaxation order of $\delta = 2$ (cf. Section 2.4) has been used throughout all the experiments and this, in general, yields a global solution to the polynomial optimization problem at hand.

Let the ground truth intrinsic calibration matrix be \mathbf{K}^0 and the estimated matrix be \mathbf{K} , where

$$\mathbf{K}^0 = \begin{bmatrix} f_1^0 & s^0 & u^0 \\ 0 & f_2^0 & v^0 \\ 0 & 0 & 1 \end{bmatrix} \text{ and } \mathbf{K} = \begin{bmatrix} f_1 & s & u \\ 0 & f_2 & v \\ 0 & 0 & 1 \end{bmatrix}.$$

Then we define the following metrics to evaluate the performance of the autocalibration algorithm:

$$\begin{aligned} \Delta f &= \frac{|\frac{1}{2}(f_1 + f_2) - \frac{1}{2}(f_1^0 + f_2^0)|}{\frac{1}{2}(f_1^0 + f_2^0)} \\ \Delta r &= \max\left(\frac{r}{r^0}, \frac{r^0}{r}\right), \text{ where } r = \frac{f_1}{f_2}, r^0 = \frac{f_1^0}{f_2^0} \\ \Delta p &= \left|\frac{1}{2}(|u| + |v|) - \frac{1}{2}(|u^0| + |v^0|)\right| \\ \Delta s &= |s - s^0|. \end{aligned} \quad (16)$$

For a calibration matrix of the form $\text{diag}(f, f, 1)$, the ideal values of these metrics are $\Delta f = 0$, $\Delta r = 1$, $\Delta p = 0$ and $\Delta s = 0$. An additional quantity of interest is the chirality check, that is, the (percentage) number of cameras that follow the chirality constraints.

For the synthetic experiments, there are 12 cameras and 15 points. All the points lie in the cube $[-1, 1]^3$ of side length 2, centered at the origin and the nominal distance of the cameras from the origin is 2. For the case of constant focal length, ground truth intrinsic calibration matrix for each camera is $\mathbf{K}^0 = \text{diag}(1, 1, 1)$. To simulate variable focal length, the intrinsic calibration matrix is of the form $\text{diag}(f, f, 1)$, where f is allowed to attain any random value in the range $[0.05, 1.00]$. Noise with standard deviation 0.2% of the image size is added to the image coordinates prior to the projective factorization that forms the input to our algorithm. The objective function to be minimized is the one in (11). The results are tabulated in Table 1, where all the quantities reported are statistics over 100 trials.

The size of the LMI relaxations used to solve for \mathbf{Q}_∞^* is a function of the number of variables, the number of constraints and the maximum degree of the polynomial occurring in the objective function and constraints. For the optimization problem (11) that we solve in this paper, none of these quantities depend on the number of views. Thus the time complexity of our algorithm is essentially constant with respect to the number of views.

Note that these experiments were conducted with a relatively few number of cameras and points, which sometimes

causes the geometry to become ill-conditioned. If that happens, the algorithm can break down due to numerical instabilities, which manifest themselves as constraint violations (such as the few chirality violations above). Another way to detect this is by checking the rank of the moment matrix in the LMI relaxation which should have rank one for globally optimal solutions.

A smaller set of experiments was performed using the objective function in (13). This was found to perform marginally better than the objective function in (11) for experiments conducted with three views. But the algorithm already becomes very computationally intensive for just three views and impossible to solve within reasonable memory limits for four or more views. The gains in terms of accuracy of solution achieved by using this theoretically more correct objective function are not enough to justify the enormous computational expense, so we will restrict ourselves to the objective function (11) henceforth.

6. Experiments with real data

In [16], several image sequences for a variety of scenes are obtained with a hand-held camcorder. The number of images varies from 3 to 125. The images themselves are quite noisy and although acquired with a constant zoom setting, auto-focus effects cause focal length to vary across the sequence. The resulting projective factorizations were upgraded to metric using a variety of algorithms and their results visually compared. The results are rated on a scale of 0 (severely distorted) to 5 (very good metric reconstruction) according to the qualitative criteria listed in [16].

We evaluate the metric reconstructions obtained by our algorithm for 25 of these sequences using the same qualitative criteria as in [16]. These comparisons are tabulated in Table 2. Some sample scene reconstructions are depicted in Figure 1.¹

Method A is the method of [16] where a quasi-affine reconstruction is obtained after untwisting the cameras and a non-linear local optimization method is used to minimize an appropriate objective function which specifies some requirements on the intrinsic parameters. Method B is the technique in [3]. Method C is the algorithm in [6] which uses the full set of chirality constraints to obtain an estimate of the plane at infinity. The method of [18] is used to obtain the reconstructions in Method D by minimizing a cost function based on the absolute quadric starting from a linear initialization. Method E is a modified version of [8]. More details on the individual methods A-E can be found in the references above or [16].

Method F is the method described in this paper where we impose the requirement that the estimated dual quadric be

¹The VRML reconstructions for these and other sequences are available at <http://vision.ucsd.edu/quadric>.

Quantity	Fixed focal length		Variable focal length	
	Without chirality	With chirality	Without chirality	With chirality
Δf	0.0086	0.0360	0.0069	0.0328
Δr	1.0041	1.0315	1.0027	1.0272
Δp	0.0073	0.0096	0.0055	0.0098
Δs	0.0051	0.0092	0.0033	0.0050
chirality	N/A	0.9867	N/A	0.9942
failed	2	2	4	1

Table 1. Performance on synthetic data. The numbers are mean values over 100 trials for the performance metrics defined in (16). For fixed focal length, $\mathbf{K}^0 = \text{diag}(1, 1, 1)$. For variable focal length, $\mathbf{K}^0 = \text{diag}(f, f, 1)$, $f \in [0.05, 1.00]$. 0.2% noise is added in image coordinates. The last row indicates the number of experiments that failed due to numerical issues and were excluded from the results .

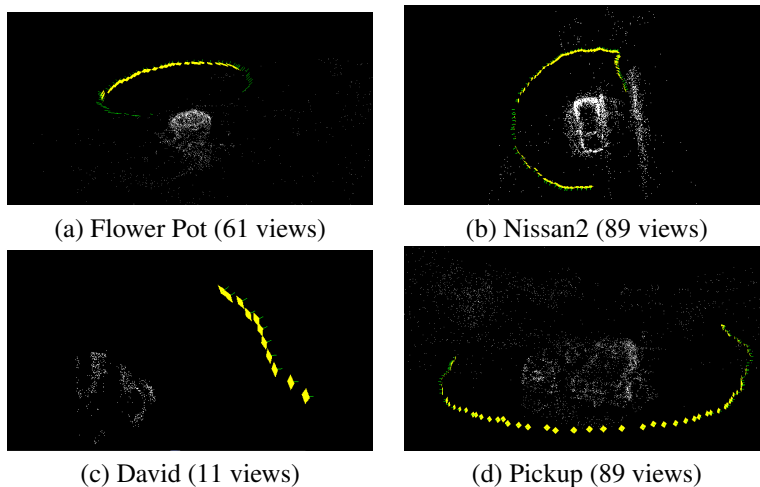


Figure 1. Metric reconstructions for four real sequences. The points are plotted in white, the image planes in yellow and the optical axes are plotted in green.

positive semidefinite and rank deficient. The optimization problem is given by (10) and the objective function used is (11). Method G is the same algorithm, but now chirality is imposed with respect to the set of camera centers.

As a general observation, the algorithm performs well for a larger number of cameras. The sequence “Basemnt” has forward camera motion, while the sequence “Drunk” has rotational camera motion. Both these cases are ill-conditioned for our algorithm and thus, the recovered structure is projectively distorted. Besides these sequences, unless the algorithm breaks down due to any numerical instability, the metric reconstructions obtained are at par with other state of the art algorithms.

Note that all our experiments are performed without imposing any bounds on the camera parameters. If one explicitly enforces prior information, for instance, that aspect ratio typically lies between 0.25 and 3, the reconstruction quality in even the scenes such as “Basemnt” which are ill-conditioned can be improved. Further, the reconstructions

evaluated above are the raw output of our algorithm. Due to the fact that interior point solvers only solve the LMI relaxation up to an ϵ tolerance, in practice a subsequent bundle adjustment step can be used to further refine the solution.

7. Conclusion

Autocalibration is a mature research topic and yet none of the previously existing approaches are capable of handling many of the hard, non-convex constraints that should be imposed according to the theory of autocalibration. In this paper we present a solution that guarantees a theoretically correct estimate of the dual absolute quadric. Experiments show that the resulting algorithm is scalable, stable and robust, with performance comparable to the state-of-the-art methods in the field.

Acknowledgments

Manmohan Chandraker and David Kriegman were supported by NSF EIA-0303622. Sameer Agarwal was funded

Dataset	Views	A	B	C	D	E	F	G
Basemnt	9	4	4	4	4	4	2	2
House	9	5	3	5	5	5	5	5
David	11	5	3	5	2	5	5	5
ClockB	13	4	3	4	3	4	4	4
Frode2	15	5	1	5	5	5	5	5
Nissan1	17	5	3	5	5	5	5	5
Stove	19	5	1	5	5	5	5	5
Drunk	21	4	1	4	4	4	*	*
SceneSw	23	5	2	5	3	5	5	5
Wine	28	4	2	4	4	4	4	4
CorrPill	35	5	1	5	5	5	5	5
FlowPt2	43	5	2	5	5	5	5	5
Contai1	57	5	2	-	5	5	5	5
Nissan3	59	5	3	5	0	5	5	5
FlowPt1	61	5	1	5	5	5	5	5
Contai2	65	5	1	-	1	5	5	5
FlowPt3	83	5	1	5	5	5	5	5
StatWk2	85	5	3	5	5	5	5	5
Nissan2	89	5	1	-	5	5	5	5
Pickup	89	5	1	-	4	5	5	5
Bicycles	103	5	5	5	1	5	5	5
GirlsSt2	105	5	1	-	2	5	5	5
StatWk1	107	5	1	5	5	5	5	5
Volvo	117	5	1	-	1	5	5	5
SwBre2	125	5	1	4	5	5	5	5

Table 2. Performance on real data. The results are rated on a scale of 0 (severely distorted) to 5 (very good metric reconstruction) according to the qualitative criteria listed in [16]. A “*” denotes a case where the algorithm encountered numerical errors due to ill-conditioning. Methods A-E are described in [16], [3], [6], [18], [8] respectively. Method F is the algorithm described in this paper. Method G is the same algorithm, but now chirality is imposed with respect to the set of camera centers.

by NSF EIA-0321235, University of Washington Animation Research Labs, Washington Research Foundation, Adobe and Microsoft. Fredrik Kahl was supported by the European Commission’s Sixth Framework Programme (SMERobot grant 011838) and by the Swedish Research Council grants 2004-4579 and 2005-3230.

References

[1] S. Agarwal, M. Chandraker, F. Kahl, S. Belongie, and D. Kriegman. Practical global optimization for multiview geometry. In *ECCV*, pages 592–605, Graz, Austria, 2006.

[2] M. Agrawal. On automatic determination of varying focal lengths using semidefinite programming. In *ICIP*, 2004.

[3] P. Beardsley, A. Zisserman, and D. Murray. Sequential updating of projective and affine structure from motion. *IJCV*, 23(3):235–259, 1997.

[4] O. Faugeras, Q.-T. Luong, and S. Maybank. Camera self-calibration: Theory and experiments. In *ECCV*, pages 321–334. Springer-Verlag, 1992.

[5] A. Fusiello, A. Benedetti, M. Farenzena, and A. Busti. Globally convergent autocalibration using interval analysis. *PAMI*, 26(12):1633–1638, 2004.

[6] R. Hartley. Euclidean reconstruction from uncalibrated views. *Applications of Invariance in Computer Vision*, 825:237–256, 1994.

[7] R. I. Hartley. Chirality. *IJCV*, 26(1):41–61, 1998.

[8] R. I. Hartley, E. Hayman, L. de Agapito, and I. Reid. Camera calibration and the search for infinity. In *ICCV*, pages 510–517, Kerkyra, Greece, 1999.

[9] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2004. Second Edition.

[10] D. Henrion and J. B. Lasserre. GloptiPoly: Global optimization over polynomials with Matlab and SeDuMi. *ACM Trans. Math. Soft.*, 29(2):165–194, 2003.

[11] A. Heyden and K. Åström. Euclidean reconstruction from constant intrinsic parameters. In *ICPR*, volume I, pages 339–343, Vienna, Austria, 1996.

[12] A. Heyden and K. Åström. Minimal conditions on intrinsic parameters for Euclidean reconstruction. In *Asian Conf. Computer Vision*, volume II, pages 169–176, 1998.

[13] F. Kahl and D. Henrion. Globally optimal estimates for geometric reconstruction problems. In *ICCV*, pages 978–985, Beijing, China, 2005.

[14] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. Optimization*, 11:796–817, 2001.

[15] R. Manning and C. Dyer. Metric self calibration from screw-transform manifolds. In *CVPR*, volume I, pages 590–597, 2001.

[16] D. Nistér. Untwisting a projective reconstruction. *IJCV*, 60(2):165–183, 2004.

[17] M. Pollefeys, R. Koch, and L. Van Gool. Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters. In *ICCV*, pages 90–95, Mumbai, India, 1998.

[18] M. Pollefeys, R. Koch, and L. Van Gool. Self-calibration and metric reconstruction in spite of varying and unknown internal camera parameters. *IJCV*, 32(1):7–25, 1999.

[19] M. Pollefeys, L. Van Gool, and M. Oosterlinck. The modulus constraint: A new constraint for self-calibration. In *ICPR*, volume I, pages 349–353, Vienna, Austria, 1996.

[20] F. Schaffalitzky. Direct solution of modulus constraints. In *Proc. Indian Conf. on Computer Vision, Graphics and Image Processing*, pages 314–321, 2000.

[21] N. Shor. *Nondifferentiable Optimization and Polynomial Problems*. Kluwer Academic Publishers, 1998.

[22] P. Sturm. A case against kruppa’s equations for camera self-calibration. *PAMI*, 22(10):1199–1204, 2000.

[23] B. Triggs. Autocalibration and the absolute quadric. In *CVPR*, pages 609–614, 1997.