

Beyond Gröbner Bases: Basis Selection for Minimal Solvers

Viktor Larsson
Lund University
Lund, Sweden

viktor1@maths.lth.se

Magnus Oskarsson
Lund University
Lund, Sweden

magnuso@maths.lth.se

Kalle Åström
Lund University
Lund, Sweden

kalle@maths.lth.se

Alge Wallis
University of Bath
Bath, United Kingdom

d.wallis@bath.ac.uk

Zuzana Kukelova
FEE

kukelzuz@fel.cvut.cz

Tomas Pajdla
CIIRC

pajdla@cvut.cz

Czech Technical University in Prague, Czech Republic

Abstract

Many computer vision applications require robust estimation of the underlying geometry, in terms of camera motion and 3D structure of the scene. These robust methods often rely on running minimal solvers in a RANSAC framework. In this paper we show how we can make polynomial solvers based on the action matrix method faster, by careful selection of the monomial bases. These monomial bases have traditionally been based on a Gröbner basis for the polynomial ideal. Here we describe how we can enumerate all such bases in an efficient way. We also show that going beyond Gröbner bases leads to more efficient solvers in many cases. We present a novel basis sampling scheme that we evaluate on a number of problems.

1. Introduction

In this paper, we describe a method for automatically building very fast minimal solvers. This is a core problem in many computer vision applications, *e.g.* 3D reconstruction [42, 43], visual odometry [36, 2] and visual localization [38, 48].

During the last years we have seen a large increase in the variety of available technical platforms, such as mobile devices, UAVs and drones. These are often equipped with widely different capabilities, in terms of sensors, cameras and computing power. In many applications, *e.g.* autonomous navigation, augmented reality or robotics, a core computer vision task is to make robust estimates of the surrounding geometry and motion of the device, based on image data [51, 11, 48, 49, 39] or other sensor data [24, 41, 3, 23]. Given that these tasks often need to be performed fast, based on unreliable data containing mismatches, and on devices with limited processing power, effi-

cient implementations of robust estimation schemes such as RANSAC is paramount. At the core of these algorithms we have so-called minimal solvers that, based on a small data sample, estimate a model that can be evaluated on a larger data set to find a consistent inlier set. Since this has to be performed many times [14], we need these minimal solvers to be fast. The image formation process naturally leads to geometric problems that can be formulated in terms of multivariate polynomial equation systems. In many devices we have additional sensor measurements, *e.g.* gyroscope and accelerometer data, and these measurements should be incorporated in the estimation process [1, 18, 48]. We can also have different types of camera models and calibration knowledge [51, 10, 50, 34, 31, 49, 17, 39]. All the aspects described above, naturally lead to the need for tools for constructing fast solvers of polynomial equations. In addition, the variety of platforms leads to the need for these methods to be, to a large extent, automatic.

Many state-of-the-art polynomial solvers are based on Gröbner bases and the action-matrix method, and there are now powerful tools available for the automatic generation of such solvers [26, 29, 30]. In this paper we target a specific part of this pipeline, namely the choice of monomial basis in the quotient ring. We will show how careful selection of the monomial bases can give significant speed-up in the resulting solvers. Previously, little attention has been paid to the choice of basis to gain speed in polynomial solvers, and usually a Gröbner basis is used to select the monomial basis. We will in the paper describe how we can test *all possible* Gröbner bases. We will further show that going beyond Gröbner bases leads to faster solvers in a number of cases.

Specifically, our contributions in this paper are:

- Minimizing elimination template size by enumerating all possible Gröbner bases.

- A heuristic method of sampling monomial bases that goes beyond Gröbner bases.
- State-of-the-art performance on a number of geometric estimation and calibration problems in terms of speed.

1.1. Background and Related Work

We use nomenclature and basic concepts from Cox *et al.* [12]. Let $X = (x_1, x_2, \dots, x_n)$ denote our indeterminates and let $\mathbb{C}[X]$ denote the set of all polynomials in X with coefficients in \mathbb{C} . We are interested in solving systems of polynomial equations,

$$\begin{cases} f_1(x_1, \dots, x_n) = 0, \\ \vdots \\ f_m(x_1, \dots, x_n) = 0. \end{cases} \quad (1)$$

The ideal $I = \langle f_1, \dots, f_m \rangle$ is the set of polynomial combinations of our generators f_1, \dots, f_m . Each polynomial $f \in I$ then vanishes on the solutions to our equation system (1). Using the ideal I we can define the quotient ring $\mathbb{C}[X]/I$ which is the set of equivalence classes over $\mathbb{C}[X]$ defined by the relation,

$$a \sim b \iff a = b \pmod{I} \iff a - b \in I. \quad (2)$$

If there are finitely many solutions to (1), then the quotient ring $\mathbb{C}[X]/I$ is a finite-dimensional vector space over \mathbb{C} .

For an ideal I there exist special sets of generators called Gröbner bases which have the nice property that the remainder after division is unique. For a Gröbner basis $\{g_1, \dots, g_m\}$ we can define the standard monomials, which is the set of monomials not divisible by the leading term of any g_k . This set of monomials is a linear basis for the quotient ring $\mathbb{C}[X]/I$.

1.2. Solving Systems of Polynomial Equations

To solve systems of polynomial equations, the most common approach in Computer Vision is the action matrix method [13, 47, 44, 28]. The goal of the action matrix method is to transform the, in general very difficult, problem of finding the solutions to an equivalent eigenvalue/eigenvector problem which we can solve numerically:

$$\alpha \begin{bmatrix} b_1 \\ \vdots \\ b_K \end{bmatrix} - \begin{bmatrix} & & \\ & M & \\ & & \end{bmatrix} \begin{bmatrix} b_1 \\ \vdots \\ b_K \end{bmatrix} = 0 \pmod{I}, \quad (3)$$

where $\alpha, b_k \in \mathbb{C}[X]$. Here α is the so-called *action variable*. To briefly motivate this, consider the quotient ring $\mathbb{C}[X]/I$ of an ideal I generated by polynomials with a finite number of solutions K . Let $b_1, \dots, b_K \in \mathbb{C}[X]$ be monomials forming a basis of $\mathbb{C}[X]/I$. Then, the remainder of the product αb_i after division by I can be written as

a linear combination of the basis $b_1, \dots, b_K \in \mathbb{C}[X]$ [13]. In matrix form we then get (3). The existence of the matrix $M \in \mathbb{C}^{K \times K}$ can be guaranteed by choosing b_1, \dots, b_K as a basis for the quotient ring $\mathbb{C}[X]/I$.

In practice, to recover the action matrix M for a particular instance a so-called *elimination template* is typically used. The elimination template is an expanded set of equations (constructed by multiplying the original equations with different monomials) in which we can linearly express the polynomials in the eigenvalue problem. This reduces the problem of finding the action matrix to solving a linear system.

It is generally difficult to find the smallest elimination template for a given problem. In [26] an iterative method was presented for constructing these elimination templates directly from the problem equations. This work was recently extended by Larsson *et al.* in [29], where a non-iterative method was proposed.

When creating polynomial solvers in Computer Vision, the basis for the quotient space $\{b_1, \dots, b_K\}$ is typically chosen as the standard monomials from the Gröbner basis w.r.t. the monomial ordering GRevLex (this is e.g. done in [26, 29]). However, this is an arbitrary choice and the methods work for any basis. In this paper we focus on the problem of selecting this basis with the aim of reducing the size of the elimination template. We show that for some problems there are better choices that yield significantly faster solvers.

Given a monomial basis it is still a difficult problem to find the smallest elimination template. In this work we use the automatic generator from [29] to construct the templates, but this method is not guaranteed to find the optimal template. The results in this paper are w.r.t. this particular template construction method. However, any other method for constructing the template (such as the one from [26]) could be substituted.

Our approach can be used to optimize other characteristics of the solvers as well, such as accuracy or stability. However, for practical purposes these are typically secondary to runtime as long as they are sufficiently good. For example if you are estimating the focal length it usually does not matter whether the errors are 10^{-6} or 10^{-16} .

In [9], Byröd *et al.* presented different methods for choosing the basis during runtime. However in their setting the size of the template was fixed, and the online basis selection was done solely to improve the numerics of the solver. In [21] the authors further improved stability by carefully selecting the so-called *permissible* monomials (*i.e.* the set from which the basis is chosen from in [9]).

2. Exhaustive Search over Gröbner Bases

For an ideal I , the (reduced) Gröbner basis depends on the monomial ordering chosen in the polynomial ring $\mathbb{C}[X]$.




Figure 1. The Gröbner fan of the ideal $I = \langle x + y^2 - 1, xy - 1 \rangle$ consists of three two-dimensional cones. For each cone, there is exactly one reduced Gröbner basis of I . All monomial orderings generated by all weight vectors from one cone give the same reduced Gröbner basis of I . Hence, there are exactly three different reduced Gröbner bases for I over all possible different monomial orderings.

Different orderings can yield different Gröbner bases, and thus different sets of standard monomials. For polynomial solvers in Computer Vision, the most popular ordering is GRevLex [12], since it has been empirically observed to typically give small elimination templates [45, 44, 26, 29]. It has also been noted in computational algebraic geometry and cryptography [40] that graded orderings [12] (*i.e.* archimedean [37]) often lead to faster Gröbner basis computations compared to, *e.g.* lexicographical orderings [12]. However, there exist examples where this is not the case. Hence, this suggests to investigate the efficiency of Gröbner basis (and hence action matrix) construction w.r.t. all possible different monomial orderings.

2.1. Gröbner Fans

While there are (uncountably) infinitely many different monomial orderings, Mora and Robbiano [33] showed that for a given ideal I there are only finitely many different reduced Gröbner bases [16]. To present this theory is beyond the scope of this paper, but we will try to describe the main ideas and relate how this can be used in our problem setting. The set of all reduced Gröbner bases of an ideal can be computed [16, 20, 19] using the *Gröbner fan* of the ideal [33, 46]. The Gröbner fan of an ideal was defined by Mora and Robbiano in 1988 [33]. It is a finite fan of polyhedral cones indexing the distinct initial monomial ideals with respect to monomial orderings or, equivalently, indexing the reduced Gröbner bases of the ideal. See [33, 46, 16] for the full account of the theory.

Here we will illustrate it on a simple example computed using the software package Gfan [19, 20]. Consider the polynomial system $I = \langle x + y^2 - 1, xy - 1 \rangle$. Figure 1 shows the Gröbner fan of I together with the corresponding reduced Gröbner bases. It consists of three two-dimensional

cones. For each cone, there is exactly one reduced Gröbner basis of I , giving in total three different reduced Gröbner bases. To connect the different reduced Gröbner bases to the fans in Figure 1, consider the exponent vectors $[a, b]^T$ that correspond to monomials $x^a y^b$, *e.g.* $[2, 3]^T$ represents $x^2 y^3$. Now, for every monomial ordering $<$ on $\mathbb{C}[x, y]$ one can find a (set of) real non-negative (weight) vectors $w \in \mathbb{R}^2$ such that if $x^a y^b < x^c y^d$, then $[a, b] \cdot w \leq [c, d] \cdot w$. In this way, every ordering is connected to a set of its (compatible) real weight vectors. Finally, for a fixed I , the union of all the sets of weight vectors corresponding to all monomial orderings producing the same reduced Gröbner basis is a full (here two) dimensional cone in \mathbb{R}^2 . There are only finitely many such cones for a fixed I . In our situation, there are three two-dimensional cones, see Figure 1.

2.2. Building Minimal Solvers using Gröbner Fans

In [29], the state-of-the-art automatic generator for polynomial solvers has been presented and evaluated on a large test-bed of polynomial equation systems from geometric Computer Vision. Even though some of the problem formulations from [29] are no longer state-of-the-art for their respective problem, they still serve as a good benchmark set to test our methods. For each of these problems we tried to compute the Gröbner fan, aborting the computations if they lasted more than 12 hours. Using the automatic generator we then construct a polynomial solver for each of the reduced Gröbner bases found. Table 1 shows some problems where we were able to find a smaller elimination template compared to using the GRevLex basis. Note that the number of Gröbner bases can increase very quickly and it is not always tractable to compute the complete Gröbner fan for larger problems. For the six point relative pose with shared radial distortion problem we ran the Gröbner fan computation for a week before aborting the computation.

Figure 2 shows a histogram of the different template sizes for the solvers constructed from the Gröbner fan for the P4Pfr formulation from Bujnak *et al.* [7]. We can see that many of the found bases give very large templates. To avoid these uninteresting bases, as well as the long runtimes for computing the Gröbner fan, we propose to use a guided random sampling approach in the next section.

3. Beyond Gröbner Bases

In the previous section we computed all reduced Gröbner bases for a problem and used these to select quotient ring bases. However, it is not necessary to select a standard monomial basis that comes from a Gröbner basis for some monomial ordering, since any spanning and linearly independent set will do. In this section we instead consider bases which do not come as standard monomial bases from any Gröbner basis, and show that for some problems we can find even smaller elimination templates.




Figure 2. Template size (rows) for the Gröbner fan bases for the P4Pfr formulation from Bujnak *et al.* [7].

3.1. Random Sampling for Basis Selection

Once you drop the Gröbner basis constraint you have infinitely many choices for monomial bases, so it is no longer possible to do any exhaustive search. Even if we restrict ourselves to monomials below some fixed degree, the combinatorial explosion of choices often makes it intractable to try them all.

Instead we propose a random sampling approach. The sampling is guided by several heuristics based on empirical observations. We will motivate our choices later on, but we will start by describing our proposed algorithm. The heuristics we use for basis selection are:

- H1. We try to have as many of the basis monomials and reducible monomials (*i.e.* αb_k) appearing in the original equations as possible.
- H2. We try to minimize the degree in some of the unknowns. This usually helps when the variables occur in an unbalanced way in the equations. E.g. if our problem is parameterized using a quaternion (for rotation) and a focal length, we have seen that it is typically good to try to minimize the degree of the focal length.
- H3. We try to select a connected block of monomials.

To generate the initial set of monomials that we sample from we use the following strategy: We start with the monomials occurring in the equations. If these do not contain any basis (see Section 3.2) we multiply with all first degree monomials that occur in the equations and add these. If they still do not contain any basis, we again multiply with all the second degree monomials and so on (in some special cases we need to add some extra low-degree monomials to get an independent set). We denote these monomials by \mathcal{M} and the monomials that occur in the original equations by $\mathcal{E} \subset \mathcal{M}$.

Now, to sample a basis we start by randomly choosing a binary weight vector $\omega = \{0, 1\}^n$. This represents the direction we want to minimize in H2. For each monomial $m \in \mathcal{M}$ we assign a weight $w_d(m)$ penalizing the weighted

degree using ω . So, e.g. if $\omega = (0, 1, 1)$, the monomial $m = xyz^2$ would have the weighted degree $0 + 1 + 2 = 3$. Next we select an action variable α . It is chosen uniformly in the direction which is minimized by ω . So, in the previous example we would have chosen either y or z . If ω is all zero we choose uniformly from all variables. Note that this α is used only for guiding the random sampling. When we construct the solvers we try every variable as action.

The basis is then sampled iteratively, with one monomial added at a time. Given a partial basis $\mathcal{B} \subset \mathcal{M}$ we select the next monomial to add as follows:

1. Find monomials $\mathcal{M}_{\mathcal{B}} \subset \mathcal{M}$ that are linearly independent from the partial basis \mathcal{B} (see Section 3.2)
2. For each monomial $m \in \mathcal{M}_{\mathcal{B}}$ compute a weight

$$w(m) = \mathbb{I}(m \in \mathcal{E}) + \mathbb{I}(\alpha m \in \mathcal{E} \cup \mathcal{B}) + w_d(m) + \epsilon \quad (4)$$
 where ϵ is a small number.
3. Find the neighboring monomials of \mathcal{B} in $\mathcal{M}_{\mathcal{B}}$.
4. Sample proportionally to $w(m)$ from the neighboring monomials. (If there are no neighboring monomials in $\mathcal{M}_{\mathcal{B}}$, sample instead from all of $\mathcal{M}_{\mathcal{B}}$).

These steps are iterated until we have a complete basis.

3.2. Checking Linear Independence

When we sample basis elements, we need to be able to quickly determine if a set of monomials are linearly independent in the quotient ring $\mathbb{C}[X]/I$ (or typically $\mathbb{Z}_p[X]/I$ since we do most of our calculations in \mathbb{Z}_p to speed up computations and avoid round-off errors).

We start by computing any (reduced) Gröbner basis for the ideal and find the standard monomials $\{b_1, b_2, \dots, b_K\}$ for this basis. Then, since these monomials form a basis for the quotient ring, we write each $m \in \mathcal{M}$ as

$$m = \sum_k c_k b_k \pmod{I}, \quad (5)$$

by simply dividing with the Gröbner basis. This associates vector $c = (c_1, c_2, \dots, c_K)$ to each monomial in \mathcal{M} . To check if a set of monomials is linearly independent in the quotient ring, we can now equivalently check if the corresponding vectors are independent in \mathbb{C}^K (or \mathbb{Z}_p^K) by performing the standard Gaussian elimination.

3.3. Building Minimal Solvers with Sampled Bases

We applied our random sampling strategy in an experiment similar to the one in Section 2.2. For each problem, we randomly sampled 100 bases and constructed the corresponding solvers. Some results are shown in Table 1. We can see that using our sampling strategy we can find

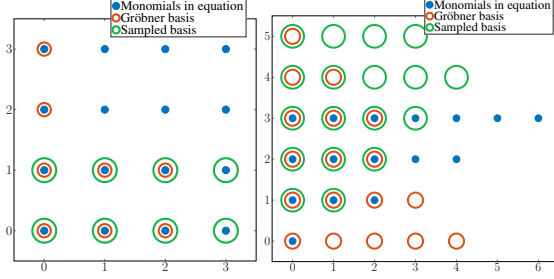


Figure 3. The figure shows the basis monomials for two example problems, namely 8pt rel. pose $F+\lambda$ (left) and 3pt image stitching $f\lambda+R+f\lambda$ (right). Both these problems have two variables, and for both these problems the proposed basis sampling scheme gives significantly smaller template compared to the Gröbner basis variants.

smaller elimination templates for some problems. Note that for some problems the best basis did not come from any Gröbner basis. We were also able to find smaller solvers for problems where the Gröbner fan computation took too long and was aborted.

In Figure 3 we show the best resulting monomial bases from our sampling scheme, for two cases. The monomials are here represented by their corresponding exponents as vectors. We also show the standard GRevLex bases, that give significantly larger templates for these two problems. In general the GRevLex ordering will lead to a basis that has a low total degree. We have found that, for some problems, if it is possible to keep the maximum degree low in one variable, even if the total degree becomes larger, this is beneficial for the template size. Figure 3 left shows an example of this. Another important aspect that we have seen, is that we should choose, if possible, monomials within the original equations, as these are available directly. In Figure 3 right, all the monomials occurring in the original equations are shown as blue dots. In this case the sampled basis better aligns with the structure of the monomials in the equations compared to GRevLex.

3.4. Experiment: Heuristic vs. Uniform Sampling

In this section we show a comparison of our heuristic with sampling basis monomials uniformly. We compare three different approaches: (i) our heuristic sampling from the monomials in \mathcal{M} (as defined in Section 3.1), (ii) uniformly sampling from \mathcal{M} , and (iii) uniformly sampling from all monomials of the same degree as those in \mathcal{M} . Figure 4 shows the distribution of the template sizes (number of rows) for 1,000 random samples for the P4Pfr formulation from Bujnak *et al.* [7]. We can see that our sampling heuristic and the strategy for selecting \mathcal{M} both give significant improvements.




Figure 4. Template size (rows) for 1,000 randomly sampled bases for the P4Pfr formulation from Bujnak *et al.* [7].

4. Panoramic Stitching $f\lambda + R + f\lambda$

We will now show how our method can be used to construct fast solvers for stitching images from cameras with radial distortion and where the focal length is unknown. This problem was formulated and solved using Gröbner-basis techniques in [8]. In [35] a technique for numerically optimizing the size of the elimination template was presented, and a new faster solver with a template of size 54×77 was constructed. In [29] a slightly faster solver was presented, based on a template of size 48×66 . We will follow the derivations in [5] and [8] when we construct our solver for two-view stitching using three point correspondences. We will additionally show that we can use the exact same solver to solve the minimal problem of three-view stitching using two point correspondences.

4.1. Two View Image Stitching

We assume that we have a camera undergoing some unknown rotation R , taking two images of a number of unknown 3D points \mathbf{X}_i . We denote the points in the two images with \mathbf{u}_i and \mathbf{u}'_i respectively. We will describe how we handle the radial distortion later, and will assume that we only need to handle the unknown focal length f just now. The projection equations can then be formulated as

$$\gamma_i \mathbf{u}_i = K \mathbf{X}_i, \quad \gamma'_i \mathbf{u}'_i = KR \mathbf{X}_i, \quad (6)$$

where γ_i and γ'_i are the depths, and $K = \text{diag}(f, f, 1)$. We can remove the dependence of γ_i , γ'_i and R by solving for \mathbf{X}_i and taking scalar products, giving the constraints

$$\frac{\langle K^{-1} \mathbf{u}_j, K^{-1} \mathbf{u}_k \rangle^2}{|K^{-1} \mathbf{u}_j|^2 |K^{-1} \mathbf{u}_k|^2} = \frac{\langle \mathbf{X}_j, \mathbf{X}_k \rangle^2}{|\mathbf{X}_j|^2 |\mathbf{X}_k|^2} = \frac{\langle K^{-1} \mathbf{u}'_j, K^{-1} \mathbf{u}'_k \rangle^2}{|K^{-1} \mathbf{u}'_j|^2 |K^{-1} \mathbf{u}'_k|^2}, \quad (7)$$

for two points j and k . Cross-multiplying with denominators will give polynomials in the unknown f . We will now add radial distortion to our problem, and model it using Fitzgibbon's division model [15] so that for the radially distorted image coordinates \mathbf{x}_i we have $\mathbf{u}_i \sim \mathbf{x}_i + \lambda \mathbf{z}_i$, where $\mathbf{z}_i = [0 \quad 0 \quad x_i^2 + y_i^2]^T$, and λ is the radial distortion parameter. Inserting this into (7) gives us our final constraints

Problem	Author	Original	[29]	GFan+ [29]	(#GB)	Heuristic+[29]
Rel. pose F+ λ 8pt	Kuang <i>et al.</i> [25]	12 \times 24	11 \times 20	11 \times 20	(10)	7 \times 16
Rel. pose E+f 6pt	Bujnak <i>et al.</i> [6]	21 \times 30	21 \times 30	11 \times 20	(66)	11 \times 20
Rel. pose f+E+f 6pt	Kukelova <i>et al.</i> [26]	31 \times 46	31 \times 50	31 \times 50	(218)	21 \times 40
Rel. pose E+ λ 6pt	Kuang <i>et al.</i> [25]	48 \times 70	34 \times 60	34 \times 60	(846)	14 \times 40
Stitching f λ +R+f λ 3pt	Naroditsky <i>et al.</i> [35]	54 \times 77	48 \times 66	48 \times 66	(26)	18 \times 36
Abs. Pose P4Pfr	Bujnak <i>et al.</i> [7]	136 \times 152	140 \times 156	54 \times 70	(1745)	54 \times 70
Rel. pose λ +E+ λ 6pt	Kukelova <i>et al.</i> [26]	238 \times 290	149 \times 205	-	?	53 \times 115
Rel. pose λ_1 +F+ λ_2 9pt	Kukelova <i>et al.</i> [26]	179 \times 203	165 \times 200	84 \times 117	(6896)	84 \times 117
Rel. pose E+f λ 7pt	Kuang <i>et al.</i> [25]	200 \times 231	181 \times 200	69 \times 90	(3190)	69 \times 90
Rel. pose E+f λ 7pt (elim. λ)	-	-	52 \times 71	37 \times 56	(332)	24 \times 43
Rel. pose E+f λ 7pt (elim. f λ)	Kukelova <i>et al.</i> [27]	51 \times 70	51 \times 70	51 \times 70	(3416)	51 \times 70
Abs. pose quivers	Kuang <i>et al.</i> [22]	372 \times 386	216 \times 258	-	?	81 \times 119
Rel. pose E angle+4pt	Li <i>et al.</i> [32]	270 \times 290	266 \times 329	-	?	183 \times 249
Abs. pose refractive P5P	Haner <i>et al.</i> [17]	280 \times 399	240 \times 324	157 \times 246	(8659)	240 \times 324

Table 1. Size of the elimination templates for some minimal problems. For the relative pose problems unknown radial distortion is denoted with λ and unknown focal length with f , and the position describes which camera it refers to. The table shows the original template size from the author, the template size found using the method from [29] (GRevLex basis), the template size from doing an exhaustive search over Gröbner bases (Section 2.2) and the random sampling approach (Section 3.1). Missing entries are when the Gröbner fan computation took longer than 12 hours.

Author	Execution time (ms)
Proposed	0.16
Larsson <i>et al.</i> [29]	0.38
Byröd <i>et al.</i> [8]	0.89

Table 2. Timing of three point stitching with unknown focal length and radial distortion, using Matlab implementations running on a standard desktop computer.

in the unknown λ and f . Using only two points will only give us one equation so we need at least three point correspondences (this actually gives three constraints, so it is slightly over-determined, but we only use two of the equations).

We have run both the exhaustive Gröbner basis selection and our proposed basis sampling scheme, see Table 1. The Gröbner bases do not give any improvement over the state-of-the-art solver but our sampling gives a significantly smaller template of size 18 \times 36. Table 2 shows a comparison of runtimes.

4.2. Three View Image Stitching

The constraints (7) only compare pairs of images, using two point correspondences. So if we, instead of having three point correspondences in two views, have two point correspondences in three views, we get the same type of constraints, namely

$$\frac{\langle K^{-1}\mathbf{u}_1, K^{-1}\mathbf{u}_2 \rangle^2}{|K^{-1}\mathbf{u}_1|^2 |K^{-1}\mathbf{u}_2|^2} = \frac{\langle K^{-1}\mathbf{u}'_1, K^{-1}\mathbf{u}'_2 \rangle^2}{|K^{-1}\mathbf{u}'_1|^2 |K^{-1}\mathbf{u}'_2|^2}, \quad (8)$$

and

$$\frac{\langle K^{-1}\mathbf{u}''_1, K^{-1}\mathbf{u}''_2 \rangle^2}{|K^{-1}\mathbf{u}''_1|^2 |K^{-1}\mathbf{u}''_2|^2} = \frac{\langle K^{-1}\mathbf{u}'_1, K^{-1}\mathbf{u}'_2 \rangle^2}{|K^{-1}\mathbf{u}'_1|^2 |K^{-1}\mathbf{u}'_2|^2}, \quad (9)$$

where double primes are used for image three. We can hence use the exact same solver to solve this case. In this case we have a true minimal case, since we only get two constraints on f and λ .

4.3. Evaluation

We have implemented our solver in MATLAB, where all image coordinate input and manipulation were done using mex-compiled C++ routines. In order to have a fair comparison of our method with [29] and [8], we modified their code so that the corresponding image coordinate manipulations also were done using mex-compiled code. The timing comparison is shown in Table 2, and one can see a clear speed-up. The solvers were run on a standard desktop computer. In order to check the numerical stability of our solver, we generated synthetic data, and evaluated the equation residuals. The results can be seen in Figure 5. In order to see how well our method works in practice, we did an automatic panoramic stitching of two images, with a fish-eye lens and unknown focal length, shown to the left in Figure 6. We then ran our solver in a standard RANSAC framework, with tentative correspondences based on SURF features and descriptors. The results can be seen to the right in Figure 6. Here the panoramic image was done without any blending in order to show the correctness of the stitching. The transformation used was based on the best RANSAC solution from our solver based on only three point correspondences, without any further bundle adjustment.




Figure 5. The figure shows histograms of equation residuals for 10,000 examples of the 3 pt stitching problem.

5. Relative Pose $E + f\lambda$

As another example we consider the relative pose problem where the calibration and distortion parameter are known for only one of the two cameras. The goal is to find a fundamental matrix F and distortion parameter λ that satisfy the epipolar constraints

$$[\hat{x}_i, \hat{y}_i, 1] F [x_i, y_i, 1 + \lambda(x_i^2 + y_i^2)]^T = 0, \quad (10)$$

as well as a focal length f , that makes

$$E = F \text{diag}(f, f, 1) \quad (11)$$

an essential matrix. The problem is minimal with seven point correspondences and has 19 solutions. The first solver was presented by Kuang *et al.* [25] and was recently improved by Kukulova *et al.* [27].

5.1. Formulation of Kuang *et al.*

Now we give a brief overview of the formulation used in Kuang *et al.* [25]. The scale of the fundamental matrix is fixed by setting $f_{33} = 1$, and the epipolar constraints yield seven equations in the monomials

$$\{\lambda f_{13}, \lambda f_{23}, \lambda, f_{11}, f_{12}, f_{13}, f_{21}, f_{22}, f_{23}, f_{31}, f_{32}, 1\}. \quad (12)$$

Using the first six equations, Kuang *et al.* linearly eliminate the first two columns of the fundamental matrix¹

$$[f_{11}, f_{12}, f_{21}, f_{22}, f_{32}, f_{33}]^T = G [\lambda f_{13}, \lambda f_{23}, \lambda, f_{13}, f_{23}, 1]^T \quad (13)$$

where $G \in \mathbb{R}^{6 \times 6}$. Finally, the last equation expresses the monomial λf_{13} as a quadratic function $h(\lambda, f_{13}, f_{22})$, which gives the additional equation

$$\lambda f_{13} - h(\lambda, f_{13}, f_{22}) = 0. \quad (14)$$

Parametrizing the inverse focal length w , the essential matrix is given by $E = F \text{diag}(1, 1, w)$, and it must satisfy the equations

$$2EE^T E - \text{tr}(EE^T)E = 0, \quad \det(F) = 0. \quad (15)$$

¹Note that here we have the focal length and distortion on the right side of F , while it was on the left in [25].

This gives 11 equations in unknowns w, λ, f_{13} and f_{23} . Using these equations, Kuang *et al.* [25] constructed a polynomial solver with a template of size 200×231 .

Computing the Gröbner fan, we found that there are 3190 different reduced Gröbner bases for this problem. Constructing solvers for all of these bases, we found a solver of size 69×90 . Applying the random approach in Section 3.1, we did not find any better solver. While this solver is significantly smaller than the original solver from Kuang *et al.* [25] (200×231), it is still slightly larger than the state-of-the-art solver from Kukulova *et al.* [27] (51×70).

5.2. Formulation of Kukulova *et al.*

In [27] the authors present another formulation for this problem based on computing elimination ideals to eliminate both the radial distortion and focal length. Since the radial distortion makes the epipolar constraints non-linear they first employ a lifting technique to remove the non-linearity. They introduce new variables y_1, y_2 and y_3 and construct an extended fundamental matrix as in [4],

$$\hat{F} = \begin{bmatrix} f_{11} & f_{12} & f_{13} & y_1 \\ f_{21} & f_{22} & f_{23} & y_2 \\ f_{31} & f_{32} & f_{33} & y_3 \end{bmatrix}, \quad (16)$$

together with the equations $y_i = \lambda f_{i3}$. Now the epipolar constraints are linear constraints on \hat{F} ,

$$[\hat{x}_i, \hat{y}_i, 1] \hat{F} [x_i, y_i, 1, x_i^2 + y_i^2]^T = 0. \quad (17)$$

Using the (now) linear constraints on \hat{F} they parametrize it using four unknowns. Finally using the elimination ideal trick they eliminate both the focal length and radial distortion parameter to get new polynomial constraints on the elements on \hat{F} . Using these new equations, they were able to construct a solver with a template size 51×70 .

We computed the Gröbner fan for this parametrization and found that there are 3416 reduced Gröbner basis. Among these we found no solver better than the GRevLex solver built by Kukulova *et al.* We also performed the random sampling approach without finding any improvement. This matches our intuition that for equation systems where the unknowns are balanced in the monomials, GRevLex performs very well.

5.3. Our Approach

Empirically we have seen that our basis selection approach works best when the monomials appear in some unbalanced way in the equations. In the parametrization from Kukulova *et al.* [27], the only unknowns are the nullspace parameters from the linear equations.

To get more imbalanced equations, we propose another formulation which is a combination of the two previous approaches. In particular, we use the elimination ideal trick



Figure 6. Stitching of two images with large radial distortion using on our three-point solver in a standard RANSAC framework. The resulting panorama (right) is based on the best RANSAC three-point solution without any additional non-linear refinement.

to eliminate the focal length, but keep the radial distortion parameter as an unknown. This avoids the extra unknowns introduced by the lifting in (16). Using similar linear eliminations as Kuang *et al.* [25], the fundamental matrix is expressed in λ , f_{13} and f_{23} . Then, instead of directly parametrizing the focal length and adding the essential matrix constraints (15), we add the eliminated constraints for one-sided focal length from [27] which only depend on the elements of the fundamental matrix. Together with the constraint from (14), we get five equations in only three unknowns. Computing the Gröbner fan, we find 332 different reduced Gröbner bases. The best solver was of size 37×56 . Finally using the random sampling approach we find a solver with an elimination template of size 24×43 .

5.4. Evaluation

We performed a synthetic experiment to evaluate the numerical stability of the new solver. We generated 1,000 random (but feasible) synthetic instances. The calibration parameters were set to $f_{gt} = 10$ and $\lambda_{gt} = -0.1$. For each solver we recorded the solution with the smallest focal length error. Figure 7 shows the distribution of the \log_{10} relative focal length error $\frac{|f-f_{gt}|}{f_{gt}}$ for all 1000 instances. The numerical stability of the new solver is similar to the solver from Kuang *et al.* [25]. Note that while the stability is worse than the solver from Kukulova *et al.* [27], it is still stable enough for practical purposes. The new solver is however significantly faster with an average runtime of 1.2 ms, compared to 10 ms for the solver from Kukulova *et al.* [27] (both solvers are implemented in MATLAB). Note that this increase in speed is not only due to the smaller elimination template, but the coefficients in the template are less complex and cheaper to compute.

6. Conclusions

We have explored how basis selection can be used to make polynomial solvers based on the action matrix method faster. The concept of Gröbner fans is an efficient representation of the possible reduced Gröbner bases that arise from (infinitely many) different monomial orderings. This gives us a tool to enumerate and test all monomial bases




Figure 7. Relative error in focal length for 1,000 random instances.

that arise from different Gröbner bases. We have shown that this gives in some cases significantly smaller elimination templates, and hence much faster solvers. We have also introduced a novel sampling scheme, that optimizes some heuristic criteria that we have experimentally found to often give small templates. Our initial motivation for sampling was that the calculation and testing of all Gröbner fans in some cases takes very (or even unfeasibly) long time, but we found that going beyond Gröbner bases can yield even smaller templates. Our motivation has here been to optimize the template size but the framework could easily be modified to optimize other criteria such as numerical stability (as was done in Kuang *et al.* [21]). We have tested our method on a large number of minimal problems, and shown that we get significant speed-ups in many cases. We have also in more depth explored how our method can be used in two applications, namely panoramic stitching with unknown focal length and radial distortion and relative pose with unknown one-sided focal length and radial distortion.

7. Acknowledgements

This work is supported by the strategic research projects ELLIIT and eSENCE, Swedish Foundation for Strategic Research project "Semantic Mapping and Visual Navigation for Smart Robots" (grant no. RIT15-0038), Wallenberg Artificial Intelligence, Autonomous Systems and Software Program (WASP) funded by Knut and Alice Wallenberg Foundation, EU-H2020 project LADIO No. 731970, the European Regional Development Fund under the project IMPACT No. CZ.02.1.01/0.0/0.0/15_003/0000468, and the Czech Science Foundation Project GACR P103/12/G084.

References

- [1] C. Albl, Z. Kukelova, and T. Pajdla. Rolling shutter absolute pose problem with known vertical direction. In *Computer Vision and Pattern Recognition (CVPR)*, pages 3355–3363, 2016. **1**
- [2] H. S. Alismail, B. Browning, and M. B. Dias. Evaluating pose estimation methods for stereo visual odometry on robots. In *the 11th International Conference on Intelligent Autonomous Systems (IAS-11)*, January 2011. **1**
- [3] E. Ask, Y. Kuang, and K. Åström. A unifying approach to minimal problems in collinear and planar toea sensor network self-calibration. In *European Signal Processing Conference (EUSIPCO 2014)*, 2014. **1**
- [4] J. H. Brito, C. Zach, K. Koeser, M. Ferreira, and M. Pollefeys. One-sided radial-fundamental matrix estimation. In *British Machine Vision Conference (BMVC)*, pages 96.12–96.12. BMVA Press, 2012. **7**
- [5] M. Brown, R. I. Hartley, and D. Nistér. Minimal solutions for panoramic stitching. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1–8. IEEE, 2007. **5**
- [6] M. Bujnak, Z. Kukelova, and T. Pajdla. 3d reconstruction from image collections with a single known focal length. In *International Conference on Computer Vision (ICCV)*, pages 1803–1810. IEEE, 2009. **6**
- [7] M. Bujnak, Z. Kukelova, and T. Pajdla. New efficient solution to the absolute pose problem for camera with unknown focal length and radial distortion. In *Asian Conference on Computer Vision (ACCV)*, pages 11–24. Springer, 2010. **3, 4, 5, 6**
- [8] M. Byröd, M. Brown, and K. Åström. Minimal solutions for panoramic stitching with radial distortion. In *British Machine Vision Conference (BMVC)*. British Machine Vision Association (BMVA), 2009. **5, 6, 7**
- [9] M. Byröd, K. Josephson, and K. Åström. Fast and stable polynomial equation solving and its application to computer vision. *International Journal of Computer Vision (IJCV)*, 2009. **2**
- [10] F. Camposeco, T. Sattler, and M. Pollefeys. Minimal solvers for generalized pose and scale estimation from two rays and one point. In *European Conference on Computer Vision (ECCV)*, pages 202–218. Springer, 2016. **1**
- [11] L. Chen, Y. Zheng, B. Shi, A. Subpa-Asa, and I. Sato. A microfacet-based reflectance model for photometric stereo with highly specular surfaces. In *International Conference on Computer Vision (ICCV)*, 2017. **1**
- [12] D. Cox, J. Little, and D. O’Shea. *Ideals, Varieties and Algorithms*. Undergraduate Texts in Mathematics. Springer-Verlag, New York, NY, USA, 2007. **2, 3**
- [13] D. A. Cox, J. Little, and D. O’Shea. *Using algebraic geometry*, volume 185 of *Graduate Texts in Mathematics*. Springer-Verlag New York, 2005. **2**
- [14] M. A. Fischler and R. C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM*, 24(6):381–395, June 1981. **1**
- [15] A. W. Fitzgibbon. Simultaneous linear estimation of multiple view geometry and lens distortion. In *Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages I–I. IEEE, 2001. **5**
- [16] K. Fukuda, A. Jensen, and R. Thomas. Computing gröbner fans. *Math. Comput.*, 76(260):2189–2212, 2007. **3**
- [17] S. Haner and K. Åström. Absolute pose for cameras under flat refractive interfaces. In *Computer Vision and Pattern Recognition (CVPR)*, pages 1428–1436, 2015. **1, 6**
- [18] N. Horanyi and Z. Kato. Multiview absolute pose using 3d-2d perspective line correspondences and vertical direction. In *International Conference on Computer Vision (ICCV)*, 2017. **1**
- [19] A. N. Jensen. Gfan, a software system for Gröbner fans. <http://home.imf.au.dk/ajensen/software/gfan/gfan.html>. **3**
- [20] A. N. Jensen. A presentation of the gfan software. In *Mathematical Software - ICMS 2006, Second International Congress on Mathematical Software, Castro Urdiales, Spain, September 1-3, 2006, Proceedings*, pages 222–224, 2006. **3**
- [21] Y. Kuang and K. Åström. Numerically stable optimization of polynomial solvers for minimal problems. In *European Conference on Computer Vision (ECCV)*, 2012. **2, 8**
- [22] Y. Kuang and K. Åström. Pose estimation with unknown focal length using points, directions and lines. In *International Conference on Computer Vision (ICCV)*, pages 529–536, 2013. **6**
- [23] Y. Kuang and K. Åström. Stratified sensor network self-calibration from toea measurements. In *21st European Signal Processing Conference*, 2013. **1**
- [24] Y. Kuang, S. Burgess, A. Torstensson, and K. Åström. A complete characterization and solution to the microphone position self-calibration problem. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3875–3879. IEEE, 2013. **1**
- [25] Y. Kuang, J. E. Solem, F. Kahl, and K. Åström. Minimal solvers for relative pose with a single unknown radial distortion. In *Computer Vision and Pattern Recognition (CVPR)*, pages 33–40. IEEE, 2014. **6, 7, 8**
- [26] Z. Kukelova, M. Bujnak, and T. Pajdla. Automatic generator of minimal problem solvers. In *European Conference on Computer Vision (ECCV)*, 2008. **1, 2, 3, 6**
- [27] Z. Kukelova, J. Kileel, B. Sturm, and T. Pajdla. A clever elimination strategy for efficient minimal solvers. In *Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2017. **6, 7, 8**
- [28] Z. Kukelova and T. Pajdla. A minimal solution to radial distortion autocalibration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(12):2410–2422, 2011. **2**
- [29] V. Larsson, K. Åström, and M. Oskarsson. Efficient solvers for minimal problems by syzygy-based reduction. In *Computer Vision and Pattern Recognition (CVPR)*, 2017. **1, 2, 3, 5, 6, 7**
- [30] V. Larsson, K. Åström, and M. Oskarsson. Polynomial solvers for saturated ideals. In *International Conference on Computer Vision (ICCV)*, 2017. **1**

- [31] V. Larsson, Z. Kukulova, and Y. Zheng. Making minimal solvers for absolute pose estimation compact and robust. In *International Conference on Computer Vision (ICCV)*, 2017. [1](#)
- [32] B. Li, L. Heng, G. H. Lee, and M. Pollefeys. A 4-point algorithm for relative pose estimation of a calibrated camera with a known relative rotation angle. In *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1595–1601. IEEE, 2013. [6](#)
- [33] F. Mora and L. Robbiano. The Gröbner fan of an ideal. *Journal of Symbolic Computation*, 6(2-3):183–208, 1988. [3](#)
- [34] G. Nakano. A versatile approach for solving pnp, pnpf, and pnpfr problems. In *European Conference on Computer Vision (ECCV)*, pages 338–352. Springer, 2016. [1](#)
- [35] O. Naroditsky and K. Daniilidis. Optimizing polynomial solvers for minimal geometry problems. In *International Conference on Computer Vision (ICCV)*, 2011. [5](#), [6](#)
- [36] D. Nistér, O. Naroditsky, and J. Bergen. Visual odometry. In *Computer Vision and Pattern Recognition (CVPR)*, pages 652–659, 2004. [1](#)
- [37] L. Robbiano. Term orderings on the polynomial ring. In *Proc. of EUROCAL '85, European Conference on Computer Algebra*, 1985. [3](#)
- [38] T. Sattler, B. Leibe, and L. Kobbelt. Efficient & effective prioritized matching for large-scale image-based localization. *IEEE Trans. Pattern Anal. Mach. Intell.*, 39(9):1744–1756, 2017. [1](#)
- [39] O. Saurer, M. Pollefeys, and G. H. Lee. A minimal solution to the rolling shutter pose estimation problem. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 1328–1334. IEEE, 2015. [1](#)
- [40] A. Segers. *Algebraic Attacks from a Gröbner Basis Perspective*. PhD thesis, Eindhoven University of Technology, 2004. [3](#)
- [41] Z. Simayijiang, S. Burgess, Y. Kuang, and K. Åström. Minimal solutions for dual microphone rig self-calibration. In *22nd European Signal Processing Conference (EUSIPCO)*, pages 2260–2264, Sept 2014. [1](#)
- [42] N. Snavely, S. M. Seitz, and R. Szeliski. Photo tourism: exploring photo collections in 3D. In *ACM SIGGRAPH'06*, 2006. [1](#)
- [43] N. Snavely, S. M. Seitz, and R. Szeliski. Modeling the world from internet photo collections. *International Journal of Computer Vision (IJCV)*, 80(2):189–210, 2008. [1](#)
- [44] H. Stewénius, D. Nistér, F. Kahl, and F. Schaffalitzky. A minimal solution for relative pose with unknown focal length. *Image Vision Comput.*, 26(7):871–877, 2008. [2](#), [3](#)
- [45] H. Stewénius, F. Schaffalitzky, and D. Nister. How hard is 3-view triangulation really? In *International Conference on Computer Vision (ICCV)*, 2005. [3](#)
- [46] B. Sturmfels. *Gröbner Bases and Convex Polytopes*. University Lecture Series. American Mathematical Society, Providence, RI, USA, 1996. [3](#)
- [47] B. Sturmfels. Solving systems of polynomial equations. In *American Mathematical Society, CBMS Regional Conferences Series, No 97*, 2002. [2](#)
- [48] L. Svärm, O. Enqvist, F. Kahl, and M. Oskarsson. City-scale localization for cameras with known vertical direction. *IEEE transactions on pattern analysis and machine intelligence*, 39(7):1455–1461, 2017. [1](#)
- [49] F. Vasconcelos, D. Peebles, S. Ourselin, and D. Stoyanov. Similarity registration problems for 2d/3d ultrasound calibration. In *European Conference on Computer Vision (ECCV)*, pages 171–187. Springer, 2016. [1](#)
- [50] J. Ventura. Structure from motion on a sphere. In *European Conference on Computer Vision (ECCV)*, pages 53–68. Springer, 2016. [1](#)
- [51] J. Ventura, C. Arth, and V. Lepetit. An efficient minimal solution for multi-camera motion. In *International Conference on Computer Vision (ICCV)*, pages 747–755, 2015. [1](#)