Minimal Constraint Relaxation for Multiview Autocalibration

Norio Kosaka
National Institute of Informatics
kosakaboat@nii.ac.jp

Timothy Duff
University of Missouri - Columbia
tduff@missouri.edu

Tomas Pajdla
CIIRC, CTU in Prague
pajdla@cvut.cz

Abstract

Polynomial systems in multiview geometry are often overconstrained, and naïvely removing equations can lead to unstable or inconsistent solutions. We revisit this issue through constraint relaxation—selectively removing equations to obtain a finite set of isolated (and, ideally, well-conditioned) solutions. Focusing on the three-view Kruppa equations for autocalibration, we introduce a minimal relaxation framework for identifying subsets of constraints leading to geometrically valid solutions. Using both symbolic computation and numerical homotopy continuation methods, we enumerate and classify all possible relaxation patterns arising from our formulation of Kruppa's equations, isolating all patterns that yield finitely-many solutions. In experiments with synthetic and real images, we demonstrate empirically that specific relaxations are well-conditioned and consistently outperform multiple baselines, including a more classical Kruppa formulation. Overall, our findings establish the minimal relaxation framework as a practical tool in multiview geometry computation. The code is available at GitHub.

1. Introduction

Minimal problems in multiview geometry arise, roughly speaking, when the number of independent constraints equals the number of unknowns. More precisely, a minimal problem is specified by a polynomial system of equations which is solvable in the sense of having a finite number of solutions. These polynomial systems serve as the algebraic engine driving the geometric solutions for important visual tasks [4, 6, 9, 20, 22, 26, 27], including 3D reconstruction, augmented reality, and camera (auto)calibration.

A well–known formulation of camera autocalibration is given by the *Kruppa equations*, which relate the fundamental matrices of multiple views to the dual image of the absolute conic (DIAC) [15]. These equations express intrinsic parameters algebraically, without relying on scene geometry. However, when the Kruppa relations from three view–pairs are combined, the resulting system of 45 bilinear equations in the five DIAC parameters becomes heavily overconstrained and typically inconsistent under realistic noise. This raises the fundamental question: *for which subsets of Kruppa constraints can we expect solvability, even with image noise, while maintaining geometric validity?*

We address this question through an exhaustive symbolic–numerical analysis of the three–view Kruppa system. Our goal is to understand how constraint selection affects both algebraic minimality and numerical stability.

On the algebraic side, we enumerate all possible *relaxations*—subsets of the full system obtained by removing equations—and identify those that have a zero-dimensional solution set using symbolic and numerical methods from algebraic geometry. The analysis reveals a clear structure: among all potential relaxations, only a small fraction (roughly 10% out of 1.2 million) are, in fact, minimal. Furthermore, we show that all minimal relaxations adhere to a single combinatorial pattern, and divide into three groups determined by the generic number of (complex) solutions.

On the numerical side, we perform large-scale experiments solving various instances of these minimal relaxations with controlled image noise using numerical *homotopy continuation (HC) methods*. Specifically, we further score each minimal relaxation offline, using Jacobian condition numbers over synthetic data as a proxy for the relaxation's numerical sensitivity. We empirically find that this scoring correlates strongly with accuracy, leading us to propose a single `Global-Best` relaxation that consistently yields the most stable solutions across both synthetic and real-world experiments. Applied to image triplets, this relaxation achieves robust and accurate autocalibration, outperforming the classical Kruppa formulation and recent Branch-and-Bound variants in both precision and stability.

Contributions:

- **Theory:** We introduce a minimal relaxation framework for Kruppa's equations in three views, conduct the first exhaustive symbolic and numerical study of constraint selection for this particular formulation of the autocalibration problem, and identify all minimal relaxations.
- **Practice:** We identify a well-conditioned `Global-Best` relaxation, selected empirically by

aggregating errors across noise levels and random seeds. Compared to previous work developing HC solvers for minimal autocalibration relaxations [6], our HC solvers are much simpler and more efficient. We also find that our `Global-Best` solver demonstrates superior stability and accuracy on synthetic and real datasets when compared to previous autocalibration methods based on Kruppa's equations [24, 28].

2. Related Work

2.1. Autocalibration and the Kruppa Equations

Kruppa's equations [13, 15, 17, 21, 24] form the classical foundation of autocalibration, relating fundamental matrices to the image of the absolute conic and enabling recovery of camera intrinsics under general motion.

A direct predecessor of our approach is the work [24], which solved three-view autocalibration via a HC method applied to five out of six Kruppa equations. Later work [23, 38] considered extensions to more than three cameras, proposing nonlinear least-squares solutions. Among more recent works based on Kruppa's equations, [28] notably introduced a globally-optimal inlier maximization approach based on branch-and-bound (BnB), combining algebraic-variety sampling with local refinement to explore feasible solutions in both DIAC and PaI space.

Beyond Kruppa's equations, autocalibration methods have been developed using various other algebraic constraints, such as those based on the dual absolute quadric (DAQ) and the modulus constraint. We refer to [15, Ch. 19] for a thorough exposition of early work based on such constraints and how it relates to the DIAC used in our approach. A more recent DAQ-based method appears in [25] which, like ours, assumes a constant intrinsics matrix $\mathbf{K}$. This approach builds on the well-known minimal problem of projective reconstruction from six point-triples; after solving this problem, the projective reconstruction may be upgraded by solving a certain set of polynomial equations. Since the DIAC is simply the conic outline of the DAQ (see e.g. [15, eq. 19.3]), DAQ-based formulations can generally be *decomposed* into solving a DIAC-based formulation and recovering $\mathbf{K}$. Algebraically, the map $\omega^*(\mathbf{K}) := \mathbf{K}\mathbf{K}^\top$ is generally four-to-one due to sign ambiguities; this symmetry was recently exploited in [6], which developed HC autocalibration solvers based on depth constraints.

2.2. Overconstrained Polynomial Systems

Solving overconstrained systems of polynomial equations may generally be approached either by (nonlinear) least-squares or exact algebraic relaxations. Least-squares methods, too prevalent to review here, are mostly dependent on some form of initial guess, whereas the exact methods commonly used in vision are more-or-less synonymous with minimal solvers. Despite their differences, least-squares and exact methods are symbiotically linked through heuristic sampling procedures, which aim to construct an adequate least-squares initialization via iterative minimal solving and outlier estimation. Even the smaller subsystems obtained from heuristic sampling, however, generally need not be solvable with noise. This motivates a more broad overview of algebraic relaxation techniques.

Heuristic sampling. Robust estimators such as RANSAC [12], MLESAC, and MSAC [37] handle overconstrained systems by randomly sampling subsets of equations that form solvable minimal systems. These methods explore the landscape of constraint selection as a stochastic process, exploring only a small fraction of the combinatorial space of subsystems. Among such estimators, PROSAC [5] further biases the sampling process by exploiting a prior ranking of correspondences, favoring higher-quality constraints early in the sampling process.

Algebraic relaxations. In this paper, every minimal relaxation is obtained by taking a subset of equations from some overconstrained systems. However, the precise notion of a minimal relaxation we adopt from [6] is more permissive. For example, this notion allows for square systems obtained by taking general linear combinations of the overconstrained equations, a procedure sometimes known as *squaring-up* (see e.g. [2, §2.5].) An even more permissive notion of minimal relaxation might allow other constructions, such as introducing auxiliary *slack variables* (see e.g. [32].) Squaring-up and slack variables are of fundamental importance in frameworks such as *numerical algebraic geometry*, which provides a toolkit for studying the solution sets of arbitrary polynomial systems, which cannot easily be studied by simply dropping equations. However, a minimal relaxation is far from arbitrary; as we recall in Section 4, it defines a variety of problem-solution pairs, which is typically irreducible and equipped with a projection onto the problem space which is generically finite-to-one. In this context, dropping equations is often preferable to squaring up or introducing slack variables, as the latter approaches yield systems that may obscure the original geometric structure by destroying sparsity or overwhelming the solver with new variables unrelated to the original problem.

Our approach. To overcome the limitations mentioned above, we exhaustively probe a natural space of minimal relaxations of Kruppa's equations for three cameras with fixed intrinsics. This yields a *complete catalog* of zero-dimensional, well-conditioned subsystems. Our relaxations of Kruppa's equations are polynomial systems with only 32, 24, or 18 solutions, and are thus much cheaper to solve than the minimal relaxations of [6], which have large root counts (for example, 2985 for the general case of $\mathbf{K}$ with all parameters unknown.) This complete catalog enables further numerical experimentation, leading ultimately to our pro-

posed `Global-Best` minimal solver.

3. Problem Setup

This section formalizes the autocalibration problem using Kruppa equations over three views, leading to an overconstrained polynomial system. We introduce the notion of *minimal relaxation* to reduce this system to a solvable form.

3.1. Autocalibration with Kruppa equations

Autocalibration aims to estimate a camera's intrinsic parameters from multiple uncalibrated views using only image correspondences. Among various formulations, the Kruppa equations offer an algebraic approach by relating the epipolar geometry of a view-pair to the dual image of the absolute conic (DIAC). Unlike conventional treatments (SM: A)[1] that rely on the singular value decomposition (SVD) of the fundamental matrix [15], we adopt a general-form Kruppa system that treats the fundamental matrix and epipoles algebraically, with only mild structural assumptions.

For a single view-pair, the Kruppa constraint can be written as $\mathbf{F}\omega^*\mathbf{F}^\top \sim [\mathbf{e}']_\times \omega^* [\mathbf{e}']_\times$ (*equality up to scale*), where $\mathbf{F} \in \mathbb{R}^{3\times3}$ is the fundamental matrix, $\mathbf{e}' \in \mathbb{R}^3$ the epipole with $\mathbf{F}^\top\mathbf{e}' = 0$, the notation $[\bullet]_\times$ indicates a 3×3 skew-symmetric matrix representing the cross product , and $\omega^* \in \mathbb{R}^{3\times3}$ is the symmetric DIAC, parametrized as

$$\omega^* = \mathbf{K}\mathbf{K}^\top = \begin{bmatrix} w_1 & w_2 & w_3 \\ w_2 & w_4 & w_5 \\ w_3 & w_5 & 1 \end{bmatrix}, \text{ where } \mathbf{K} = \begin{bmatrix} f & s & u \\ 0 & g & v \\ 0 & 0 & 1 \end{bmatrix}.$$

The intrinsic matrix $\mathbf{K}$ contains parameters (f, g, u, v, s), corresponding to focal lengths, principal point, and skew. Here, we have fixed the scale by setting ${\omega^*}_{3,3} = 1$.

For a three-view minimal problem, we consider three fundamental matrices $\{\mathbf{F}_1, \mathbf{F}_2, \mathbf{F}_3\}$ and three corresponding epipoles $\{\mathbf{e}'_1, \mathbf{e}'_2, \mathbf{e}'_3\}$, one for each view-pair. Applying the Kruppa constraint to each view yields:

$$\begin{aligned} \mathbf{F}_1\omega^*\mathbf{F}_1^\top &\sim [\mathbf{e}'_1]_\times\omega^*[\mathbf{e}'_1]_\times \leftrightarrow \mathbf{A}_1 \sim \mathbf{B}_1, \\ \mathbf{F}_2\omega^*\mathbf{F}_2^\top &\sim [\mathbf{e}'_2]_\times\omega^*[\mathbf{e}'_2]_\times \leftrightarrow \mathbf{A}_2 \sim \mathbf{B}_2, \\ \underbrace{\mathbf{F}_3\omega^*\mathbf{F}_3^\top}_{\mathbf{A}\in\mathbb{R}^{3\times3}} &\sim \underbrace{[\mathbf{e}'_3]_\times\omega^*[\mathbf{e}'_3]_\times}_{\mathbf{B}\in\mathbb{R}^{3\times3}} \leftrightarrow \mathbf{A}_3 \sim \mathbf{B}_3. \end{aligned}$$

The symmetric matrices $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3$, and $\mathbf{B}_1, \mathbf{B}_2, \mathbf{B}_3$ can be reshaped into six 6×1 vectors and stacked column-wise to form three 6×2 matrices $\mathbf{C}_1, \mathbf{C}_2, \mathbf{C}_3$, namely

$$\mathbf{C}_i = \begin{bmatrix} A_{i,1} & B_{i,1} \\ A_{i,2} & B_{i,2} \\ \vdots & \vdots \\ A_{i,6} & B_{i,6} \end{bmatrix}_{6\times2}. \tag{1}$$

[1] By "SM: A.N" we refer to Section N in the Supplementary Material.

To enforce equality up to scale, we impose the vanishing of all 2×2 minors of $\mathbf{C}_i$. Let

$$G_i = \left\{ f_i^{(p,q)} = A_{i,p}B_{i,q} - A_{i,q}B_{i,p} \mid 1 \le p < q \le 6 \right\} \tag{2}$$

denote the set of these $15 = \binom{6}{2}$ minors for view-pair i. Stacking all constraints across three view-pairs yields $15 \times 3 = 45$ polynomial equations:

$$\mathcal{G} = \{G_1, G_2, G_3\} = \{f_1, \ldots, f_{45}\}. \tag{3}$$

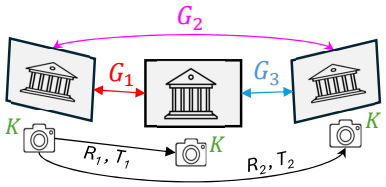

Figure 1. Illustrating the setup of 3-view Kruppa systems.

The polynomial system $\mathcal{G}$ compactly encodes the algebraic constraints of minimal autocalibration from three views. With 45 equations in 5 unknowns $\mathbf{x} = (w_1, \ldots, w_5)$, it is little surprise that this system is overconstrained. We note that, from any exact, physically-meaningful solution $\mathbf{x}$, Cholesky decomposition of the associated 3×3 matrix ω^* will give us the intrinsics matrix $\mathbf{K}$.

4. How to Relax

We next introduce the minimal relaxation framework on an example of an overconstrained camera autocalibration problem using Kruppa equations. In short, we formalize how solvable subsystems arise by selectively removing constraints from the full set of Kruppa equations.

Geometric estimation problems in vision (including autocalibration) can be viewed as algebraic maps between spaces of parameters and data. Let $\pi : \mathcal{Z}_{P,X} \to P$ denote the projection from an (irreducible) *problem–solution variety* $\mathcal{Z}_{P,X} \subseteq P \times X$, whose points $(\mathbf{p}, \mathbf{x})$ pair a problem instance $\mathbf{p} \in P$ with a corresponding solution $\mathbf{x} \in X$. The projection $\pi(\mathbf{p}, \mathbf{x}) = \mathbf{p}$ takes as input a pair consisting of a problem instance $\mathbf{p}$ and one of its solutions $\mathbf{x}$, and simply outputs the problem instance $\mathbf{p}$.

Definition 1. [9] A problem is said to be *minimal* if:

1. The projection π is **dominant**, meaning its image is dense in the problem space P (in the Zariski topology).
2. For a **generic** problem $\mathbf{p} \in P$, the fiber $\pi^{-1}(\mathbf{p})$ (i.e. the solution set) is finite and nonempty.

The *degree* of the problem is the cardinality of a generic fiber, $\deg(\pi) := |\pi^{-1}(\mathbf{p})|$, i.e. the number of complex solutions corresponding to a generic problem instance $\mathbf{p}$.

Intuitively, a minimal problem is one where the number of independent constraints matches the number of unknowns, so that a generic instance admits finitely many valid solutions. In practice, minimality can be verified locally: if $\mathbf{p}_0 \in P$ is generic and $(\mathbf{p}_0, \mathbf{x}_0) \in X$ is contained in a neighborhood locally defined by the system $\mathcal{F}(\mathbf{p}, \mathbf{x}) = 0$, the problem is minimal precisely when

$$\operatorname{rank}\left(\left.\frac{\partial \mathcal{F}}{\partial \mathbf{x}}\right|_{(\mathbf{p}_0, \mathbf{x}_0)} \right) = n,$$

where $n = \dim X$ is the number of unknowns [6, 9].

To recover solvability under overconstrained systems, we seek smaller subsystems $\mathcal{F} \subset \mathcal{G}$ that remain algebraically valid while satisfying the rank condition above.

Definition 2. (c.f. [6]) Given an overconstrained polynomial system $\mathcal{G}$ vanishing on a problem-solution variety $\mathcal{Z}_{P,X}$, a subset $\mathcal{F} \subset \mathcal{G}$ is called a *minimal relaxation* if it defines a minimal problem-solution variety in some neighborhood of a generic point $(\mathbf{p}, \mathbf{x}) \in \mathcal{Z}_{P,X}$.

Due to genericity in this definition, the minimality of $\mathcal{F}$ persists under generic parameter variations, making the relaxation well-defined rather than instance-specific.

In the context of Kruppa's equations, the problem space $P \ni (\mathbf{F}_1, \mathbf{F}_2, \mathbf{F}_3)$ is the variety of *equi-uncalibrated* compatible fundamental matrix triples (denoted Y_Δ in [10]), the unknown space is $X = \mathbb{C}^5 \ni \omega^*$, and the irreducible variety $\mathcal{Z}_{P,X}$ is the image of a rational joint-image map [9]. In the following sections, we analyze minimal relaxations $\mathcal{F}$ obtained by dropping equations from the system $\mathcal{G}$ defined in equation (3), which locally defines the variety $\mathcal{Z}_{P,X}$.

4.1. Algebraic analysis of constraint selection

We now study which subsets of Kruppa constraints $\mathcal{G}$ have finite solution sets—a necessary condition for any minimal relaxation $\mathcal{F}$. Letting $V(\bullet)$ denote the solution set of any given polynomial system, we generally have the inequality $\dim V(\mathcal{F}) \geq \dim V(\mathcal{G})$ for any subset $\mathcal{F} \subset \mathcal{G}$. After fixing noiseless problem data $(\mathbf{F}_1, \mathbf{F}_2, \mathbf{F}_3)$, we are looking for relaxations where this inequality is an equality, namely $\dim V(\mathcal{F}) = \dim V(\mathcal{G}) = 0$.

Dimension and Degree for Kruppa Relaxations. Our analysis of minimal relaxations for the Kruppa system of equation (3) used symbolic and numerical methods provided by the computer algebra system `Macaulay2` [14]. We first used symbolic Gröbner basis computation to compute the dimensions of the polynomial ideals generated by candidate relaxations $\mathcal{F}$ obtained from noiseless problem data. We further verified the zero-dimensional cases by a Jacobian rank test, to ensure that no spurious positive-dimensional components influenced the results.

We also determined the *degree* of each minimal relaxation $\mathcal{F}$—that is, the number of points in the solution set $|V(\mathcal{F})|$, using both symbolic Gröbner basis calculation and the numerical monodromy method of [8]. Results based on symbolic computation, despite assuming generic data, are consistent over 50 random seeds. Further details for these experiments are provided in SM: B.

Relaxation Patterns: We consider relaxations of the general Kruppa constraints obtained by selecting 5 equations distributed across the 3 view-pairs. Since each view-pair contributes up to 5 Kruppa constraints, we explore all valid combinations of distributing the 5 equations among the 3 view-pairs. This yields all ordered triples (a, b, c) such that $a + b + c = n$ with $0 \leq a, b, c \leq 5$, resulting in distinct relaxation patterns (e.g., $(1, 1, 3)$ means we select 1, 1, and 3 equations from the first, second, and third pairs, respectively). This yields $\binom{45}{5} \approx 1.2$M combinations to study.

Results: Table 1 reports the dimensions of solution sets for all minimal relaxations of the full Kruppa system (3). Since different systems with the same relaxation pattern may still give different relaxations, we report each combination of relaxation pattern, dimension, and degree witnessed by at least one occurrence among $\binom{45}{5}$ total relaxations. When $\dim V(\mathcal{F}) > 0$, the column **Deg** reports the degree of the ideal $\langle \mathcal{F} \rangle$. For the zero-dimensional relaxations $\mathcal{F}$, we sometimes needed to saturate the ideal $\langle \mathcal{F} \rangle$ by an ideal defining spurious solutions to obtain an ideal whose degree equaled the root count supplied by monodromy HC. These occurrences comprise a subset of those counted by the three rows in Table 1 marked in **red**.

Notably, our experiments show that all relaxation patterns except $(1, 2, 2)$ consistently yield an infinite (positive-dimensional) solution set. In contrast, the $(1, 2, 2)$ relaxation pattern always yields a finite (zero-dimensional) solution set. Extending this analysis, we provide dimensions and degrees for calibration tasks under different prior assumptions (e.g. fixed aspect ratio) in SM: C.3. Overall, these results highlight that only a few relaxation types lead to algebraically minimal systems, motivating further analysis of their structure and numerical behavior.

Table 1. Dimensions and degrees of minimal relaxations for the Kruppa system (3), grouped by relaxation pattern.

Relax.	Dim	Deg	Occ.	Relax.	Dim	Deg	Occ.
(0,0,5)	3	3	2,922	(0,2,3)	1	12	27,456
	3	4	81		1	16	15,936
(0,1,4)	2	6	18,180		2	8	2,574
	2	8	2,295		2	16	1,782
(1,1,3)	1	12	64,350		3	4	27
	1	16	37,350	**(1,2,2)**	**0**	**18**	**91,250**
	2	8	675		**0**	**24**	**42,130**
					0	**32**	**4,860**
					1	16	25,920
					2	8	1,215

5. Combinatorial and numerical analyses

Table 1 shows that $\dim V(\mathcal{F}) > 0$ for most of the possible relaxations. In addition to narrowing the search for our proposed `Global-Best` relaxations, this also suggests a question: *how do such underconstrained relaxations of the system* (3) *manifest?* In Sec. 5.1, we give a precise answer to this question. Additionally, we are still left with a fairly large search space of **138,240** zero-dimensional relaxations. To find our proposed `Global-Best` relaxation, we tame this search space by exploiting permutation symmetry in Section 5.2 and empirically evaluating the Jacobian condition numbers of candidate relaxations in Section 5.3.

5.1. Characterizing Underconstrained Relaxations

To characterize the underconstrained relaxations of the Kruppa system (3), we perform symbolic calculations in a quotient of a polynomial ring with 9 indeterminate entries representing the entries of $\mathbf{F}$, modulo the prime, homogeneous, principal ideal generated by the determinant,

$$S := \mathbb{Q}[F_0, \ldots, F_8]/\langle \det \mathbf{F} \rangle.$$

Over the ring S, we define the matrices

$$\mathbf{F} = \begin{pmatrix} F_0 & F_1 & F_2 \\ F_3 & F_4 & F_5 \\ F_6 & F_7 & F_8 \end{pmatrix}, \qquad \mathbf{e}' = \begin{pmatrix} -F_4F_6 + F_3F_7 \\ F_1F_6 - F_0F_7 \\ -F_1F_3 + F_0F_4 \end{pmatrix}, \tag{4}$$

so that the symbolic relations $\det \mathbf{F} = 0$ and $\mathbf{F}^\top \mathbf{e}' = 0$ hold. Since $\langle \det \mathbf{F} \rangle \subset \mathbb{Q}[F_0, \ldots, F_8]$ is a prime ideal, we may form the field of fractions $\mathbb{K} = \mathrm{Frac}(S)$. We may then perform calculations in the polynomial ring $\mathbb{K}[\mathbf{x}]$, whose elements are polynomials in the indeterminate entries of ω^* whose coefficients belong to the field $\mathbb{K}$. Forming the matrix $\mathbf{C} = [\mathbf{a} \mid \mathbf{b}]_{6\times 2}$ as in (1), each 2×2 minor of this matrix represents an element of $\mathbb{K}[\mathbf{x}]$.

Consider now two such minors, giving nonzero elements $f_0, f_1 \in \mathbb{K}[\mathbf{x}]$. If $f_0/f_1 \in \mathbb{K}$, then these two minors are algebraically independent, which implies any relaxation containing them both is non-minimal. This simple observation gives rise to a *field-membership test* that, surprisingly, filters out all non-minimal relaxations of the Kruppa system $\mathcal{G}$. For each pair (f_0, f_1) with $f_0 \neq 0$, we compute a reduced rational function representing the quotient $f_0/f_1 \in \mathrm{Frac}(S[\omega^*])$ using `Macaulay2`. More precisely, this reduced quotient has the form $-s_0/s_1$, where the pair (s_0, s_1) encodes a *syzygy relation* $s_0f_0 + s_1f_1 = 0$. Once this reduced rational function is determined, we examine the variables appearing in its support: if no ω_i appears, the pair is declared *dependent*; otherwise, it is *independent*.

Among the 105 possible combinations, precisely nine pairs are dependent in the sense defined above. These collapse to a single intrinsic constraint after eliminating the projective variables, while all other pairs remain independent. These nine dependent pairs identified above account exactly for all underconstrained relaxations in Table 1. Specifically, when a dependent pair appears in **either** the second or the third view-pair, the resulting subsystem exhibits dimension 1, yielding 25,920 such cases. When dependent pairs occur in **both** the second and third view-pairs simultaneously, the subsystem becomes dimension 2, producing 1,215 cases. Together, these $25{,}920 + 1{,}215 = 27{,}135$ instances correspond precisely to all underconstrained relaxations reported in Table 1, confirming a one-to-one alignment between algebraic dependence and the empirically observed dimensionality.

Numerical validation of these findings in `Julia` is provided in SM: C.1.

5.2. Symmetry Reductions

Permutation equivalence arises naturally when the same relaxation pattern is realized by different choices from the three view–pairs. Relaxations differing by such permutations are algebraically equivalent up to reordering of the view–pair indices.

Specifically, a minimal relaxation with relaxation pattern $(1, 2, 2)$ selects a single constraint from the first view–pair and subsets of two from the other view-pairs. The permutations of this relaxation, which have relaxation pattern $(2, 1, 2)$ and $(2, 2, 1)$, reorder the assignment of these subsets among the three view–pairs. A related operation, the *intra-pattern swap*, fixes the view–pair with one constraint and exchanges the two subsets containing two constraints. This leaves the relaxation pattern unchanged.

Example. We illustrate both (a) a *permutation* of the view-pairs and (b) *intra-pattern swapping* of constraint subsets across view–pairs for the $(1, 2, 2)$ relaxation. Each color below denotes a fixed subset of Kruppa constraints—red for $\{1\}$, blue for $\{2, 3\}$, and green for $\{4, 5\}$. We assume that a labeling of 2×2 minors of the matrices $\mathbf{C}_i$ in (1) by the integers $1, \ldots, 15$ is provided and independent of i.

(a) *Permutation:* $(1, 2, 2) \to (2, 1, 2) \to (2, 2, 1)$

$$\{\{1\}, \{2,3\}, \{4,5\}\} \to \{\{4,5\}, \{1\}, \{2,3\}\} \to \{\{2,3\}, \{4,5\}, \{1\}\},$$

(b) *Intra-pattern swap:*

$$\{\{1\}, \{2,3\}, \{4,5\}\} \to \{\{1\}, \{4,5\}, \{2,3\}\}.$$

In each case above, the constituent subsets of constraints are permuted across view-pairs.

Algebraically, both view-pair permutations and intra-pattern swaps amount to simple relabeling of variables and parameters for any given relaxation. Thus, any two minimal relaxations related by these two symmetries effectively produce essentially the same sets of constraints. We validate this observation experimentally in SM: C.2.

5.3. Sensitivity Analysis of Minimal Relaxations

After filtering permutation–equivalent relaxations, the search space reduces from $\sim$ 1.2M to $\sim$ 84K permutation-inequivalent minimal relaxations. To prioritize which relaxations to investigate further, we measure Jacobian condition numbers associated with each relaxation and their correlations with numerical stability and accuracy.

Jacobian Condition Number. Let $\mathcal{F}(\mathbf{x}, \mathbf{p})$ denote a relaxed subsystem of Kruppa equations. Consider the Jacobian matrix $\mathbf{J_x} = \partial\mathcal{F}/\partial\mathbf{x} \in \mathbb{R}^{5\times5}$ for this relaxation, evaluated at a ground-truth solution $\mathbf{x}$ producing noiseless parameters $\mathbf{p}$. The *condition number* of this matrix may be defined as $\kappa(\mathbf{J_x}) = \frac{\sigma_{\max}(\mathbf{J_x})}{\sigma_{\min}(\mathbf{J_x})}$, where $\sigma_{\max}$ and $\sigma_{\min}$ are the largest and smallest singular values of $\mathbf{J_x}$, respectively. This condition number estimates the sensitivity of the ground-truth solution $\mathbf{x} = [\omega_1, \omega_2, \omega_3, \omega_4, \omega_5]^\top$ with respect to perturbations of system parameters $\mathbf{p}$, and is closely related to other condition numbers studied in multiview geometry [11]. In general, high condition numbers are correlated with highly sensitive solutions. Importantly, the accuracy of any numerical method that recovers the ground-truth is fundamentally limited by $\kappa(\mathbf{J_x})$. Since HC solvers in particular must ultimately solve linear systems with coefficient matrix $\mathbf{J_x}$, the condition number of this matrix is an attractive metric for further filtering minimal relaxations.

Remark: Up until this point, we have only relied on *generic properties* (in the usual sense of algebraic geometry, see e.g. [33, Ch. 4]) to filter out candidate minimal relaxations. However, we stress that the following sensitivity analysis of condition numbers is purely empirical, and *does not* measure any generic properties of minimal relaxations.

Setup. We evaluate all 84K relaxations on synthetic image data under zero Gaussian noise to isolate algebraic sensitivity from data perturbations. For each relaxation, we compute the Jacobian condition number $\kappa(\mathbf{J_x})$ at the fabricated solution and the sum of mean relative parameter errors (`sum_mpe`, see SM D.4) using the computed error-minimizing solution $\mathbf{x}$, averaged over 100 random scenes. Both quantities are reported as per-relaxation means and standard deviations. Experiments in Section 6 use normalized device coordinates (NDC) with the `PolyhedralHC` solver. Further details can be found in (SM: B).

Results. Figure 2 shows the scatter of condition number versus sum–error across all relaxations averaged over 100 random scenes. Although the plot appears sparse due to the logarithmic scale, a clear positive correlation emerges: relaxations with larger $\kappa(\mathbf{J_x})$ exhibit higher `sum_mpe`, indicating degraded numerical stability. Conversely, clusters with $\kappa(\mathbf{J_x}) < 10^5$ consistently achieve sub–10^{-8} error magnitudes, suggesting that well–conditioned Jacobians correspond to robust, accurate ground-truth solutions.

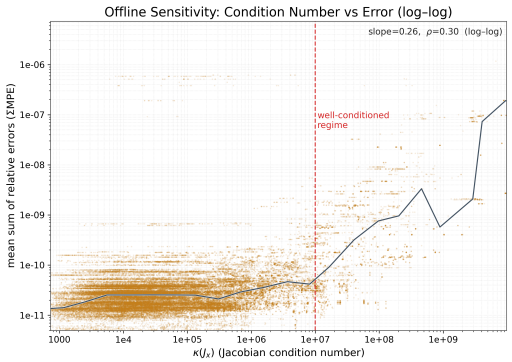

Figure 2. **Sensitivity Analysis.** Scatterplot of Jacobian condition number $\kappa(\mathbf{J_x})$ versus summed mean parameter errors (`sum_mpe`) for all 84K relaxations (log–log scale). Lower $\kappa(\mathbf{J_x})$ values correspond to more stable and accurate solutions.

6. Experiments

This section demonstrates the practical impact of the proposed minimal relaxations. We empirically assess numerical accuracy on synthetic and real-world data, showing that the selected relaxation remains robust and accurate across varying perturbations in image measurements.

Methods:

- **Baseline (`Kruppa`).** We implement the classical Kruppa formulation [24], enforcing the original Kruppa constraints via the SVD-based parametrization of the fundamental matrix (SM: A).
- **Baseline (`Kruppa-BnB`).** This is a state-of-the-art autocalibration method from [28], which employs a Branch-and-Bound (BnB) strategy with local refinement. For the main experiments (Fig. 3), we evaluate only the BnB stage of this approach without Levenberg–Marquardt. This provides a fair comparison with our proposed global initialization, which may also benefit from refinement. Further results including LM refinement are in (SM: D.1).
- **Ours (`Global-Best`).** We select the `Global-Best` relaxation *offline* from synthetic experiments (Sec. 6.1) and apply it unchanged *online* to real-image experiments (Sec. 6.2). For Fig. 3, we additionally report `Local-Best` and `Local-Worst` relaxations—the smallest and 95$^{\text{th}}$-percentile largest errors at each noise level—to illustrate variability across possible relaxations.

Both `Kruppa` and `Global-Best` were implemented using a HC package implemented in the `Julia` programming language, `HomotopyContinuation.jl` [3]. For fairer comparison between `Kruppa` and `Global-Best`, we use a polyhedral homotopy which tracks 32 paths in each case. However, we note that monodromy-initialized parameter homotopy could also be used to track 18 or 24 paths for these problems, respectively. See (SM: B.3) for more details and (SM: D.3) for runtime experiments.

Metrics: We evaluate intrinsic accuracy via normalized errors in both synthetic and real-world settings:

- **Focal length error** (Δfg): average absolute deviation of the estimated focal lengths, $\Delta f = \frac{1}{2}\left(\frac{|\hat{f}-f|}{\hat{f}} + \frac{|\hat{g}-g|}{\hat{g}}\right)$;
- **Principal point error** (Δuv): average relative error of the estimated image center, $\Delta uv = \frac{1}{2}\left(\frac{|\hat{u}-u|}{u} + \frac{|\hat{v}-v|}{v}\right)$
- **Skew error** (Δs): scaled absolute deviation of the skew term, $\Delta s = \frac{2|\hat{s}-s|}{f+g}$.

6.1. Experiment: Synthetic Images

We evaluate the robustness of Kruppa systems under varying synthetic perturbations to images of size 640×480. For each trial, the polynomial system parameters—a triplet of fundamental matrices and their associated epipoles—are generated using minimal solvers in epipolar geometry with 6-, 7-, and 8-point configurations [6]. The full experimental setup is described in (SM: B.2).

Results. Figure 3 compares the proposed `Global-Best` method with the `Kruppa` baseline under increasing pixel noise $\sigma \in [0.0, 1.0]$ across the 6-, 7-, and 8-point configurations averaged over 100 random seeds. `Global-Best` consistently yields lower mean projection error, indicating greater robustness to image noise and improved numerical stability. Averaged across all noise levels, `Global-Best` outperforms `Kruppa` in every configuration, confirming the benefit of selecting the most consistent constraint subset within each epipolar formulation. Comparable low-variance behavior across noise levels has also been reported by [6] (see their Fig. 4), suggesting inherent stability in well-conditioned minimal formulations. Further ablations on image-coordinate scaling and parameter homotopy continuation are provided in (SM: D.4). Finally, note that the branch-and-bound strategy is incompatible with the Kruppa system: the computed ω^* (DIAC) must remain positive definite for Cholesky-based recovery of $\mathbf{K}$, which is frequently violated, leading to failures with intrinsic errors on the order of thousands. For completeness, `Kruppa-BnB` results are thus included in (SM: D.1).

6.2. Experiment: Real-world images

We evaluate our offline-selected `Global-Best` relaxation on real multi-view datasets, testing its stability and accuracy under realistic camera and measurement conditions.

Datasets. We evaluate our approach on several real-world multi-view datasets commonly used in structure-from-motion and multi-view calibration, including (`SfM-eval`) the Strecha SfM benchmark [35], (`strecha08`) the CVLAB–EPFL multi-view stereo dataset [34], and (`colmap-public`) COLMAP's public dataset [31], covering various calibrated sequences from these benchmarks. Detailed scene configurations, image counts, and image resolutions are provided in (SM: B.4).

Each scene provides ground-truth 3D points, camera intrinsics, rotations, and translations for more than three calibrated views. We reconstruct each scene using `COLMAP` [30], initialized with the provided ground-truth intrinsics. The recovered 3D points paired with poses and intrinsics are used in the experiments run with `Julia`. We use the 6-point solver since it provides the most numerically stable estimation of fundamental-matrix triplets on real data; notably, Table 2 of [6] reports that the `Kruppa` variant combined with the 6-point solver achieved the best performance among the 6-, 7-, and 8-point configurations.

MSAC-based solvers. All solvers are evaluated within the MSAC framework [37]. At each iteration, a random *6-point* subset is sampled, and the corresponding 3D points are projected onto the three views using the fixed camera extrinsics to estimate a fundamental matrix triplet. This process is repeated for up to *200 iterations*, and the best result is selected as the one yielding the smallest reprojection error across all iterations. We report relative errors to quantify the deviation of estimated intrinsics from the ground truth.

Results. Table 2 summarizes the results of the real-world experiments using the 6-point solver. Overall, the proposed `Global-Best` relaxation consistently outperforms the classical `Kruppa` baseline across all datasets, with substantial gains on `SfM-eval` and `colmap-public`, and moderate yet consistent improvements on `strecha08`. A minor exception occurs on `Fountain P11`, where both methods yield similar accuracy. We note that the MSAC-guided sampling of 3D point triplets further enhances solver stability by favoring well-conditioned configurations during fundamental-matrix estimation. Real images are substantially higher in resolution than the synthetic data, resulting in proportionally larger absolute errors (SM: B.4).

7. Conclusion

We introduced a framework for selecting minimal relaxations of overconstrained polynomial systems and demonstrated its effectiveness on Kruppa's equations. Our analysis identifies the $(1, 2, 2)$ relaxation family as a stable minimal configuration with lower algebraic complexity than prior formulations [6]. Experiments show that our solver achieves accurate and robust autocalibration, outperforming conventional approaches. Code and experiments will be released.

Acknowledgements

TP was supported by OPJAK CZ.02.01.01/00/22 008/0004590 Roboprox project.

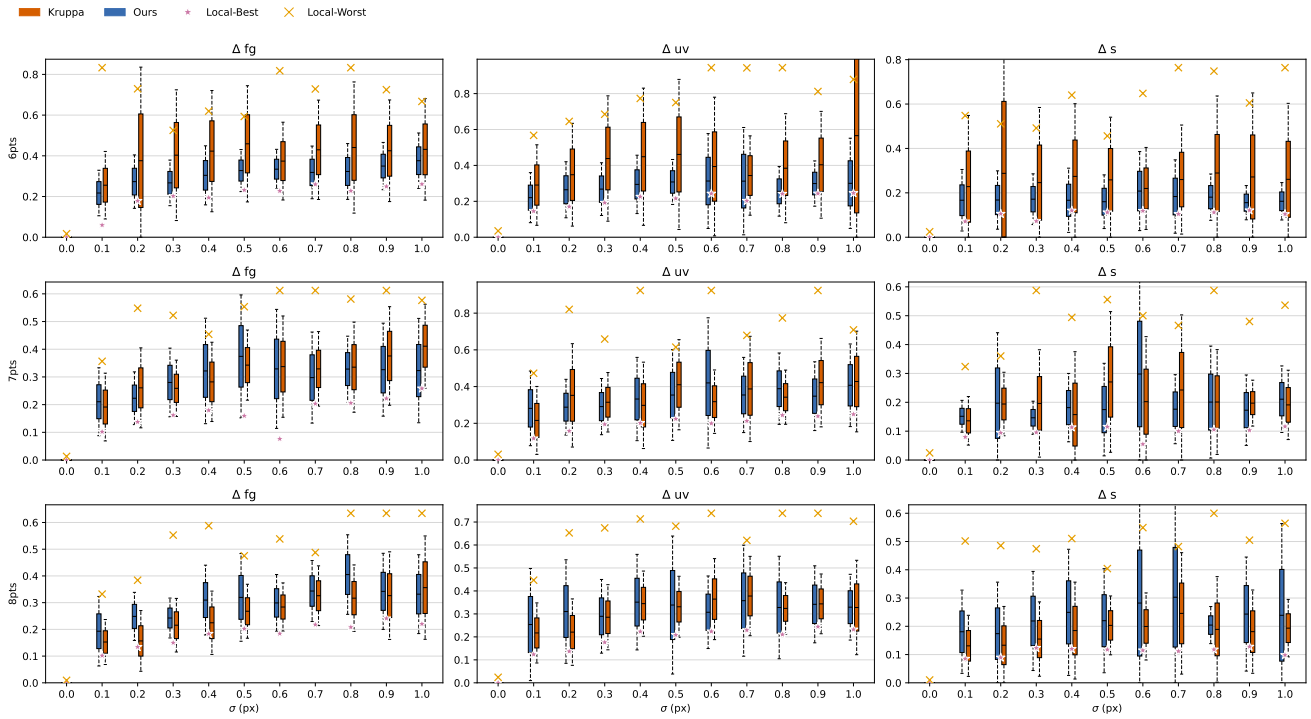

Figure 3. Synthetic comparison of solvers with the 6-, 7-, and 8-point epipolar estimations under increasing noise levels. Boxplots show the grouped errors in intrinsic parameters. Note that $\Delta_* = 0.1$ corresponds to approximately 50–80 px absolute deviation in intrinsics, indicating stable calibration even under strong noise (SM: D.2).

Table 2. Average over 2000 random scenes using the 6-point epipolar configuration. Lower is better. Metrics: Δfg, Δuv, Δs. The smallest entry in each row is highlighted in bold.

Dataset	**Scene**	**Kruppa**			**Global-Best**			**Kruppa-BnB**		
		Δfg	Δuv	Δs	Δfg	Δuv	Δs	Δfg	Δuv	Δs
SfM-eval	Castle P19	0.761	0.782	0.6151	**0.677**	**0.759**	**0.3922**	480.775	283.561	1.3765
SfM-eval	Entry P10	0.508	1.075	0.5556	**0.375**	**0.739**	**0.3485**	480.398	283.561	1.3598
SfM-eval	Fountain P11	0.649	**1.636**	**0.2918**	**0.585**	1.847	0.3896	471.568	283.561	1.3522
SfM-eval	Herz Jesus P8	1.669	2.878	0.6424	**0.699**	**1.783**	**0.3899**	434.872	283.561	1.5095
colmap-public	Gerrard Hall	0.477	0.720	**0.3132**	**0.414**	**0.517**	0.3410	684.793	279.939	1.7481
colmap-public	South Building	**0.614**	**0.511**	**0.1817**	0.699	0.654	0.2536	515.688	279.895	1.4229
strecha08	Brussel	0.576	0.601	**0.2697**	**0.536**	**0.458**	0.3004	469.599	278.658	1.3257
strecha08	Rathaus	0.689	1.623	0.0605	**0.503**	**0.976**	**0.0487**	327.455	258.245	0.1077
strecha08	Semper	0.537	0.900	0.5978	**0.446**	**0.490**	**0.3739**	489.752	282.284	1.3517
Average		0.750	1.251	0.4018	**0.565**	**0.963**	**0.3121**	458.763	279.166	1.2257

References

[1] Daniel Barath, Tekla Toth, and Levente Hajder. A minimal solution for two-view focal-length estimation using two affine correspondences. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6003–6011, 2017. 13

[2] Daniel J Bates, Paul Breiding, Tianran Chen, Jonathan D Hauenstein, Anton Leykin, and Frank Sottile. Numerical nonlinear algebra. *arXiv preprint arXiv:2302.08585*, 2023. 2

[3] Paul Breiding and Sascha Timme. Homotopycontinuation. jl: A package for homotopy continuation in julia. In *International congress on mathematical software*, pages 458–465. Springer, 2018. 6, 12

[4] Martin Byröd, Klas Josephson, and Kalle Åström. A Column-Pivoting Based Strategy for Monomial Ordering in Numerical Gröbner Basis Calculations. In *Computer Vision - ECCV 2008, Marseille, France, October 12-18, 2008, Proceedings, Part IV*, pages 130–143. Springer, 2008. 1

[5] Ondrej Chum and Jiri Matas. Matching with prosac-progressive sample consensus. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, pages 220–226. IEEE, 2005. 2

[6] Andrea Porfiri Dal Cin, Timothy Duff, Luca Magri, and Tomas Pajdla. Minimal perspective autocalibration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5064–5073, 2024. 1, 2, 4, 7, 11, 12, 13, 15, 18, 19

[7] David Cox, John Little, Donal O'shea, and Moss Sweedler. *Ideals, varieties, and algorithms*. Springer, 1997. 12

[8] Timothy Duff, Cvetelina Hill, Anders Jensen, Kisun Lee, Anton Leykin, and Jeff Sommars. Solving polynomial systems via homotopy continuation and monodromy. *IMA Journal of Numerical Analysis*, 39(3):1421–1446, 2019. 4, 13, 16

[9] Timothy Duff, Kathlen Kohn, Anton Leykin, and Tomas Pajdla. Plmp-point-line minimal problems in complete multi-view visibility. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1675–1684, 2019. 1, 3, 4

[10] Timothy Duff, Viktor Korotynskiy, Anton Leykin, and Tomas Pajdla. Multiprojective geometry of compatible triples of fundamental and essential matrices. *arXiv e-prints*, pages arXiv–2602, 2026. 4

[11] Hongyi Fan, Joe Kileel, and Benjamin Kimia. On the instability of relative pose estimation and ransac's role. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8935–8943, 2022. 6

[12] Martin A Fischler and Robert C Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6):381–395, 1981. 2

[13] Guillermo Gallego, Elias Mueggler, and Peter Sturm. Translation of" zur ermittlung eines objektes aus zwei perspektiven mit innerer orientierung" by erwin kruppa (1913). *arXiv preprint arXiv:1801.01454*, 2017. 2

[14] Daniel R. Grayson and Michael E. Stillman. Macaulay2, a software system for research in algebraic geometry. Available at http://www2.macaulay2.com. 4

[15] Richard Hartley and Andrew Zisserman. Multipleview geometry in computer vision second edition, 2004. 1, 2, 3, 11

[16] Richard I Hartley. In defense of the eight-point algorithm. *IEEE Transactions on pattern analysis and machine intelligence*, 19(6):580–593, 1997. 12

[17] Richard I. Hartley. Kruppa's equations derived from the fundamental matrix. *IEEE Transactions on pattern analysis and machine intelligence*, 19(2):133–135, 1997. 2, 11

[18] Janne Heikkila. Using sparse elimination for solving minimal problems in computer vision. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 76–84, 2017. 13

[19] Anders Heyden and Kalle Åkström. Minimal conditions on intrinsic parameters for euclidean reconstruction. In *Asian Conference on Computer Vision*, pages 169–176. Springer, 1998. 13

[20] Joe Kileel and Kathlén Kohn. Snapshot of algebraic vision. *arXiv preprint arXiv:2210.11443*, 2022. 1

[21] Erwin Kruppa. *Zur Ermittlung eines Objektes aus zwei Perspektiven mit innerer Orientierung*. Hölder, 1913. 2

[22] Viktor Larsson, Magnus Oskarsson, Kalle Åström, Alge Wallis, Zuzana Kukelova, and Tomas Pajdla. Beyond Gröbner Bases: Basis Selection for Minimal Solvers. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 3945–3954, 2018. 1

[23] Cheng Lei, Fuchao Wu, Zhanyi Hu, and Hung-Tat Tsui. A new approach to solving kruppa equations for camera self-calibration. In *2002 International conference on pattern recognition*, pages 308–311. IEEE, 2002. 2

[24] Q-T Luong and Olivier D Faugeras. Self-calibration of a moving camera from point correspondences and fundamental matrices. *International Journal of computer vision*, 22 (3):261–289, 1997. 2, 6, 11, 17

[25] Evgeniy Martyushev. A minimal six-point auto-calibration algorithm. *arXiv preprint arXiv:1307.3759*, 2013. 2

[26] Pedro Miraldo, Tiago Dias, and Srikumar Ramalingam. A Minimal Closed-Form Solution for Multi-perspective Pose Estimation using Points and Lines. In *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XVI*, pages 490–507, 2018. 1

[27] D. Nistér. An efficient solution to the five-point relative pose problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(6):756–770, 2004. 1

[28] Danda Pani Paudel and Luc Van Gool. Sampling algebraic varieties for robust camera autocalibration. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 265–281, 2018. 2, 6, 17

[29] Frederik Schaffalitzky, Andrew Zisserman, Richard I Hartley, and Philip HS Torr. A six point solution for structure and motion. In *European conference on computer vision*, pages 632–648. Springer, 2000. 12

[30] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE con-*

ference on computer vision and pattern recognition, pages 4104–4113, 2016. 7, 18

[31] Johannes L. Schönberger. Colmap dataset collection. https://demuc.de/colmap/datasets/, 2025. Accessed: 8 Nov. 2025. 7

[32] Andrew J Sommese, Jan Verschelde, and Charles W Wampler. Numerical irreducible decomposition using phcpack. In *Algebra, Geometry and Software Systems*, pages 109–129. Springer, 2003. 2

[33] Andrew J Sommese, Charles W Wampler, et al. *The Numerical solution of systems of polynomials arising in engineering and science*. World Scientific, 2005. 6, 12

[34] Christoph Strecha. Multi-view stereo dataset – cvlab, epfl. https://www.epfl.ch/labs/cvlab/data/data-strechamvs/, 2004. Accessed: 8 Nov. 2025. 7

[35] Christoph Strecha, Wolfgang Von Hansen, Luc Van Gool, Pascal Fua, and Ulrich Thoennessen. On benchmarking camera calibration and multi-view stereo for high resolution imagery. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. Ieee, 2008. 7

[36] Peter Sturm. On focal length calibration from two views. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, pages II–II. IEEE, 2001. 13

[37] Philip HS Torr and Andrew Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer vision and image understanding*, 78(1):138–156, 2000. 2, 7

[38] Cyril Zeller and Olivier Faugeras. *Camera self-calibration from video sequences: the Kruppa equations revisited*. PhD thesis, INRIA, 1996. 2