# General Self-Calibration with Varying Intrinsics*

Norio Kosaka[1,2], Timothy Duff[3], Tomas Pajdla[4], and Akihiro Sugimoto[2]

[1] Program of Informatics, SOKENDAI, Tokyo, Japan
[2] National Institute of Informatics, Tokyo, Japan
{kosakaboat,sugimoto}@nii.ac.jp
[3] University of Missouri, Columbia, MO, USA
tduff@missouri.edu
[4] CIIRC, Czech Technical University in Prague, Czechia
pajdla@cvut.cz

**Abstract.** We address multiview self-calibration from image correspondences when camera intrinsic parameters vary across views. While Kruppa equations constrain an implicit matrix encoding of the camera intrinsics (called the dual image of the absolute conic; DIAC), practical calibration priors are expressed directly in intrinsic-parameter space. We formulate such priors algebraically and map them into DIAC space, yielding explicit constraints that integrate directly with Kruppa relations. Unlike prior self-calibration methods that focus on specific varying-intrinsic regimes (e.g. shared focal length, fixed aspect, or zero skew), we provide a unified algebraic framework that handles arbitrary intrinsic priors expressed as polynomial constraints in the space of intrinsics, yielding a more flexible formulation of varying-intrinsic self-calibration. Since Kruppa systems are projective and often algebraically dependent, we algorithmically construct locally independent square subsystems via Jacobian analysis to assess solvability and algebraic complexity under varying intrinsics. Experiments demonstrate parity with a few previous classical approaches in shared-focal settings as well as shared-principal-point settings and enable stable estimation for varying-intrinsic configurations previously considered unsolved, validated on synthetic and real dynamic-intrinsic sequences. Code is available at `https://github.com/Rowing0914/Self-Calibration-Varying-Intrinsics`.

## 1 Introduction

Self-calibration is a key ingredient in unstructured multiview reconstruction: it upgrades projective geometry recovered from correspondences into a metric reconstruction by exploiting constraints on the camera intrinsics, enabling physically meaningful 3D, consistent scale-up-to-similarity, and reliable downstream tasks such as SfM/SLAM, AR, and measurement [11, 16].

While classical self-calibration has been extensively studied—ranging from bundle-adjustment-based approaches under varying focal length and principal point [17, 18], to DAQ-based metric upgrades via the dual absolute quadric [24],

* TP was supported by OPJAK CZ.02.01.01/00/22008/0004590 Roboprox project.

and Kruppa-equation formulations using the dual image of the absolute conic (DIAC) [14, 22, 23]—most analyses either assume constant intrinsics or rely implicitly on *complete-intersection assumptions* [16, 24], positing the feasibility of calibration when there are as many constraints as geometric degrees of freedom. These assumptions may not hold in genuinely varying-intrinsic regimes . Real capture pipelines frequently exhibit dynamic intrinsics due to zoom, refocus, or post-processing (e.g. electronic stabilisation), motivating formulations that explicitly address algebraic dependence in the varying-intrinsic setting.

We study Kruppa-based self-calibration under *task-dependent intrinsic priors* (e.g. zero skew, fixed aspect, shared principal point across views) when intrinsics vary across cameras. A key mismatch is that such priors are naturally expressed in the space of intrinsics ($K$-space), whereas Kruppa equations live in the space of DIACs $\omega_i^* \sim K_i K_i^\top$. To bridge this gap, we treat intrinsic priors as an algebraic set in $K$-space and map them through the map $K_i \mapsto \omega_i^*$, yielding explicit algebraic constraints in DIAC space that can be combined with Kruppa relations in a single polynomial system. Since the resulting system may be overdetermined, we use a minimal-relaxation strategy similar to [5, 20]: locally independent subsystems are sampled via Jacobian-based tests to form square polynomial systems that preserve local solvability. This enables systematic analysis of algebraic complexity while maintaining numerical tractability.

We validate the proposed formulation on both synthetic and real-world data. On synthetic scenes, we analyse algebraic root counts and study numerical stability under controlled image noise, showing that locally independent minimal relaxations yield stable intrinsic estimates. On real video sequences with frame-level ground-truth intrinsics, we evaluate both genuine focal variation and simulated principal-point shifts induced by electronic stabilisation. The method remains robust under noisy epipolar estimation, reducing to classical behaviour in the shared-intrinsic regime while accommodating genuinely varying intrinsics without modification.

*Our main contributions are:*

1. **General formulation for varying-intrinsic self-calibration.** We formulate intrinsic-task constraints directly in $K$-space and map them to DIAC space, enabling a unified Kruppa-based framework for general self-calibration with varying intrinsics under structured priors.
2. **Solving previously unsolved varying-intrinsic tasks.** We demonstrate Kruppa-based self-calibration beyond the classical shared-focal setting, including configurations with zero skew and fixed aspect ratio with a shared principal point, as well as cases with varying focal lengths but a shared principal point. Solver behaviour under pixel noise is analysed via root-count/bound analysis and validated on synthetic and real experiments.

## 2 Related Work

*Self-calibration under varying intrinsics.* Early work by Heyden and Åström [17, 18] established that Euclidean reconstruction remains feasible under continuously varying focal length and principal point, given structural priors such as zero skew or fixed aspect ratio. Pollefeys *et al.* [24] analysed which intrinsic constraints suffice to resolve projective ambiguity and introduced the Dual Absolute Quadric (DAQ) for metric upgrade, together with counting arguments relating the number of constrained degrees of freedom to the 8-parameter projective–metric gap. DAQ methods operate on a projective reconstruction and estimate a global dual absolute quadric, whereas Kruppa constraints eliminate the plane at infinity and impose conditions directly on pairwise epipolar geometry and per-view DIACs. While these works clarify solvability conditions, practical estimation typically depends on strong priors and reliable initialisation before bundle adjustment.

Kruppa-equation formulations remain central when only fundamental matrices are available. Lourakis and Deriche [22] considered varying intrinsics within a reduced Kruppa framework combined with non-linear refinement, and subsequent work improved optimisation stability [9]. We use the Kruppa form because it gives polynomial constraints directly in the DIAC variables to which our pushforward intrinsic-prior constraints are added. This does not imply that Kruppa constraints are preferable to DAQ formulations, but gives a compact setting for analysing minimal relaxations across intrinsic-prior patterns. Classical DIAC constraints for assumptions such as zero skew, square pixels, known principal point, or constant intrinsics are well known [9,11,24,26]; our focus is the systematic pushforward derivation of such constraints and the minimal-relaxation analysis of the resulting Kruppa systems. However, existing approaches either rely on heuristic initialisation (e.g. centre principal point assumptions), many-view regimes, or complete-intersection arguments. In contrast, we explicitly construct locally independent minimal subsystems and analyse their algebraic complexity, enabling principled estimation under genuinely varying intrinsic parameters.

*Two-view and specialised intrinsic solvers.* A complementary line of work studies focal length or partial intrinsic recovery from two views, leveraging properties of the fundamental or essential matrix [3, 12, 19, 27], and more recent minimal solvers using affine correspondences or elimination techniques [1, 15]. For the practically important constant-focal-length case, Bocquillon *et al.* [2] studied a specialised DAQ/projective-reconstruction approach under square-pixel and known-principal-point assumptions. These approaches are effective for restricted camera models or initialisation, but they do not directly address the multiview Kruppa/DAQ setting with broadly varying intrinsics.

## 3 Varying-Intrinsic Kruppa Self-Calibration

We now develop an algebraic formulation of self-calibration under *varying intrinsics* using Kruppa constraints. Rather than assuming a shared calibration, we retain independent DIACs per view, enabling task-dependent intrinsic priors and yielding overdetermined systems addressed via a minimal relaxation framework.

### 3.1 Kruppa constraints with varying intrinsics

Self-calibration estimates camera intrinsics directly from epipolar geometry. In the Kruppa formulation, the unknowns are the dual images of the absolute conic (DIACs), which encode the intrinsic matrices.

For each view $i = 1, \dots, m$, we allow a fully varying intrinsic matrix $K_i$ and its associated DIAC, defined projectively by $\omega_i^* \sim K_i K_i^\top$. The DIAC is a symmetric $3 \times 3$ matrix defined up to scale. We adopt the standard upper-triangular parameterisation:

$$K_i = \begin{bmatrix} f_i & s_i & u_i \\ 0 & g_i & v_i \\ 0 & 0 & 1 \end{bmatrix}, \qquad \omega_i^* = \begin{bmatrix} o_{i1} & o_{i2} & o_{i3} \\ o_{i2} & o_{i4} & o_{i5} \\ o_{i3} & o_{i5} & 1 \end{bmatrix},$$

where the projective scale is fixed by $(\omega_i^*)_{33} = 1$.

When intrinsics vary across views, the DIACs $\{\omega_i^*\}_{i=1}^m$ are independent unknowns, yielding five scalar degrees of freedom per camera after scale normalisation.

For a pair of views $(i, j)$, let $F_{ji} \in \mathbb{R}^{3\times 3}$ denote the fundamental matrix and $e'_{ji}$ the epipole in image $j$. The Kruppa constraint links the DIACs of the two views:

$$F_{ji}\, \omega_i^*\, F_{ji}^\top \;\sim\; [e'_{ji}]_\times\, \omega_j^*\, [e'_{ji}]_\times. \tag{1}$$

where $[x]_\times$ denotes the $3\times 3$ skew-symmetric matrix. We do not assume a shared DIAC across views; unlike classical constant-intrinsic Kruppa formulations that estimate a single $\omega^*$ for all cameras [11, 14, 23], we retain view-specific DIAC variables, allowing varying intrinsics and explicit task-dependent cross-view priors.

The Kruppa relation (1) is projective: both sides are symmetric $3\times 3$ matrices that must agree up to scale. For a given view pair $(i, j)$, define

$$A_{ji} := F_{ji}\, \omega_i^*\, F_{ji}^\top, \qquad B_{ji} := [e'_{ji}]_\times\, \omega_j^*\, [e'_{ji}]_\times.$$

Both $A_{ji}$ and $B_{ji}$ are symmetric. We extract their six distinct entries and stack them as 6-vectors:

$$\boldsymbol{a}_{ji} \in \mathbb{R}^6, \qquad \boldsymbol{b}_{ji} \in \mathbb{R}^6.$$

The projective equality $A_{ji} \sim B_{ji}$ is equivalent to $\boldsymbol{a}_{ji}$ and $\boldsymbol{b}_{ji}$ being collinear. We encode this collinearity by forming the $6 \times 2$ matrix

$$C_{ji} = \begin{bmatrix} a_{ji,1} & b_{ji,1} \\ a_{ji,2} & b_{ji,2} \\ \vdots & \vdots \\ a_{ji,6} & b_{ji,6} \end{bmatrix}_{6\times 2}$$

where $a_{ji,k}, b_{ji,k}$ denote the $k$-th entries of $\boldsymbol{a}_{ji}$ and $\boldsymbol{b}_{ji}$, respectively, and enforcing the vanishing of all its $2\times 2$ minors. This is algebraically equivalent to the classical ratio-based derivations of Kruppa constraints [11, 14, 22, 23]. Concretely,

$$G_{ji} = \big\{ f_{ji}^{(p,q)} = a_{ji,p}\, b_{ji,q} - a_{ji,q}\, b_{ji,p} \;\big|\; 1 \le p < q \le 6 \big\}.$$

Each pair $(i, j)$ therefore contributes $\binom{6}{2} = 15$ quadratic polynomial constraints in the DIAC variables. Let $\mathcal{V}$ denote the set of selected view pairs. For each $(j, i) \in \mathcal{V}$, let $G_{ji}$ denote the corresponding Kruppa constraint. We collect all such constraints into the family

$$\mathcal{G} \;:=\; \{\, G_{ji} \mid (j, i) \in \mathcal{V} \,\}.$$

Stacking the constraints $G_{ji}$ over all selected view pairs yields a clearly *overdetermined* polynomial system in $\omega^* = (\omega_1^*, \ldots, \omega_m^*)$ with five unknown scalars per view (after scale fixing). Task-dependent intrinsic priors (Sec. 4.1) are incorporated as additional polynomial constraints on DIAC variables, allowing all relations to be solved within a unified algebraic system.

### 3.2 Minimal Self-Calibration Problems with Relaxations

We introduce the minimal relaxation framework for overdetermined Kruppa-based self-calibration by formalising how solvable subsystems arise through selective removal of constraints from the full equation set.

Geometric estimation problems in vision (including self-calibration) 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 (e.g., image measurements) $\mathbf{p} \in P$ with a corresponding solution (e.g., DIAC) $\mathbf{x} \in X$. The projection $\pi(\mathbf{p}, \mathbf{x}) = \mathbf{p}$ maps each solution to its underlying problem instance.

**Definition 1.** *[5, 8] A problem is said to be* minimal *if:*

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

Intuitively, a minimal problem is one in which the number of independent constraints balances the number of unknowns, so that a generic instance admits finitely many solutions. In practice, minimality can be checked locally: if $\mathcal{F}(\mathbf{p}, \mathbf{x}) = 0$ defines the solution set $X$ in a neighbourhood of $(\mathbf{p}_0, \mathbf{x}_0)$, where $\mathbf{p}_0 \in P$ is generic, then the problem is minimal whenever

$$\operatorname{rank}\left( \left.\frac{\partial \mathcal{F}}{\partial \mathbf{x}}\right|_{(\mathbf{p}_0, \mathbf{x}_0)} \right) = n, \tag{2}$$

where $n = \dim X$ is the number of unknowns [5, 8].

To recover solvability under *overdetermined* systems, we seek smaller subsystems $\mathcal{F} \subset \mathcal{G}$ that remain algebraically valid while satisfying (2).

**Definition 2.** *[5] Given an overdetermined 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 near a generic point $(\mathbf{p}, \mathbf{x}) \in \mathcal{Z}_{P,X}$.*

Due to genericity [25, Ch. 4] in this definition, the minimality of $\mathcal{F}$ persists under generic parameter variations, making the relaxation well-defined rather than instance-specific (e.g., Critical Motion Sequence [26]). A **minimal relaxation** therefore consists of selecting a subset of constraints whose removal restores algebraic minimality—i.e., yields a square, locally independent system—while preserving the geometric validity of the original Kruppa formulation. As in standard self-calibration, our analysis concerns generic configurations; a complete classification of critical motion sequences and formulation-specific Kruppa degeneracies is beyond the scope of this work.

### 3.3 Sampling Minimal Relaxations

In practice, finding minimal subsystems of Kruppa system from such large pool ($\mathcal{G}$) is combinatorially prohibitive (See `#Subsys` in Tab. 1). We therefore require a principled procedure to *sample* locally independent subsystems that define minimal self-calibration problems.

*Jacobian-based independence test.* Let $\mathcal{G} = \{g_1, \ldots, g_M\}$ denote the full set of kruppa constraints across all view-pairs of interest, and let $\mathbf{x} \in X$ denote the unknowns (e.g. DIACs). We define the full Jacobian matrix

$$J := J_{\mathcal{G}}(\mathbf{p}_0, \mathbf{x}_0), \quad J(\mathbf{p}, \mathbf{x}) = \frac{\partial \mathcal{G}}{\partial \mathbf{x}} = \begin{bmatrix} \nabla_{\mathbf{x}} g_1 \\ \vdots \\ \nabla_{\mathbf{x}} g_M \end{bmatrix},$$

evaluated at a generic point $(\mathbf{p}_0, \mathbf{x}_0)$ on the problem–solution variety $\mathcal{Z}_{P,X}$. By (2), minimality requires selecting $n = \dim X$ constraints whose corresponding rows in $J$ are linearly independent.

*Constructing minimal subsets.* To find linearly independent rows in $J$, we employ a rank-revealing QR factorisation on $J^\top$ with column pivoting, yielding indices of independent rows of $J$; we retain the first $r = \text{rank}(J)$ pivot rows (with a numerical tolerance) to obtain an independent subsystem (see SM. C.2 for a comparison of the rank-revealing behaviour of QR, SVD, and LU). In this scheme, the minimal relaxation depends on the ordering of constraints $g_1, \ldots, g_M$. To diversify the selected subsystems, we randomise the row order and apply random positive row scaling of $J$, which preserves row independence but changes pivot tie-breaking. Ultimately, if $r = n$, this produces square subsystem $\mathcal{F} \subset \mathcal{G}$ with Jacobian $J_{\mathcal{F}}$ satisfying (2).

*Empirical conditioning analysis.* We instantiate this analysis on `ffuv0-uv` (Sec. 5), a representative varying-intrinsic task used in our experiments with nontrivial solution ambiguity. While $\mathrm{rank}(J_{\mathcal{F}}) = n$ ensures local minimality, we can ask more specifically for the condition number $\kappa(J_{\mathcal{F}}) = \frac{\sigma_{\max}(J_{\mathcal{F}})}{\sigma_{\min}(J_{\mathcal{F}})}$, where $\sigma_{\max}$ and $\sigma_{\min}$ denote the largest and smallest singular values of $J_{\mathcal{F}}$, respectively. This provides one natural measure of how sensitive an exact solution to a minimal problem is to noise or perturbations in the input.

To this end, we generate $N$ minimal relaxations per task on generic scenes (Sec. 4.2) and analyse the distribution of $\kappa(J_{\mathcal{F}})$ evaluated at the generic scenes. For $N \in \{10, 50, 100, 1000\}$, the mean condition numbers (mean $\pm$ std), averaged over 10 randomly sampled generic scenes, are 2.91e3 $\pm$ 2.45e3, 6.59e3 $\pm$ 1.27e4, 5.42e3$\pm$9.43e3, and 3.86e3$\pm$5.42e3, respectively. Across tasks, $\kappa(J_{\mathcal{F}})$ consistently lies within a moderate regime (typically $\lesssim 10^6$), suggesting that the sampled subsystems are generally well-conditioned. Already $N = 10$ locally independent relaxations provide a sufficiently high probability of obtaining well-conditioned subsystems. Accordingly, we fix $N = 10$ for all subsequent experiments. Further statistics are reported in SM. C.1.

## 4 Structured Intrinsic Priors in DIAC Space

This section formulates task-dependent intrinsic priors directly in DIAC space and analyses how they constrain the resulting Kruppa self-calibration system.

### 4.1 Pushing Intrinsic Priors Forward to DIAC space

For $m$ views, let $K_i$ denote the intrinsic matrix of view $i$. We encode task-dependent intrinsic priors as an algebraic set

$$P_{\mathcal{K}} \;:=\; \Big\{ K = (K_1, \ldots, K_m) \in (\mathbb{C}^5)^m \;\Big|\; F_{\mathrm{task}}(K) = 0 \Big\}, \tag{3}$$

where $F_{\mathrm{task}}$ is a collection of known *polynomial* relations in the intrinsics parameters. In practical vision settings (e.g. zero skew $s_i = 0$, square pixels $f_i = g_i$, or cross-view sharing $f_i = f_j$), intrinsic priors are typically linear in the parameters. We therefore focus on degree-1 constraints in this work. The same pushforward construction applies to higher-degree polynomial priors, at the expense of increased algebraic complexity.

Now, DIACs for each of the $m$ cameras gives us a rational map,

$$\pi_{\omega^*} : (\mathbb{C}^5)^m \dashrightarrow (\mathbb{P}^5)^m \text{ s.t. } (K_1, \ldots, K_m) \mapsto (K_1 K_1^\top, \ldots, K_m K_m^\top) = \omega^*. \tag{4}$$

(Note: Although $\pi_{\omega^*}$ is given componentwise by quadratic polynomials, it is undefined at certain points in its domain, e.g. when some $K_i$ is the zero matrix.)

Using the rational map (4), we may naturally associate to any given priors (3) an algebraic feasible set in the DIAC space,

$$P_{\mathcal{W}} \;:=\; \overline{\pi_{\omega^*}(P_{\mathcal{K}})} \;\subset\; (\mathbb{P}^5)^m, \tag{5}$$

the Zariski closure of the image of $\pi_{\omega^*}$ (see e.g. [6, Ch. 8].) We call $P_{\mathcal{W}}$ the *pushforward variety*; this is an irreducible variety whenever $P_{\mathcal{K}}$ is. Intuitively, $P_{\mathcal{W}}$ describes all $m$-tuples of DIACs $(\omega_1^*, \ldots, \omega_m^*)$ that are compatible with the chosen intrinsic prior, without explicit reference to $K_i$. We note that, even when the task priors of (3) are given by linear constraints, the equations defining the pushforward variety need not be linear due to the quadratic map (4).

Symbolic Gröbner basis computation enables us to compute *pushforward constraints*—that is, polynomial equations vanishing on the algebraic set $P_{\mathcal{W}}$. Details and examples using Macaulay2 [10] are provided in SM.B.2.

*Degrees of freedom:* The dimension $d := \dim(P_{\mathcal{W}})$ of the pushforward variety is the number of independent DIAC degrees of freedom remaining after enforcing the intrinsic prior. Each locally independent Kruppa equation (Sec. 3.3) generically reduces the DIAC solution space dimension by one. Therefore, under generic configurations, at least $d$ independent self-calibration constraints are required to obtain a *0-dimensional (minimal)* system.

*Example: Moving camera across views [14, 23].* Assume the task prior $K_1 = K_2$. In intrinsic coordinates $(f_i, g_i, u_i, v_i, s_i)$ this is the set of linear equalities

$$f_1 - f_2 = 0, \quad g_1 - g_2 = 0, \quad u_1 - u_2 = 0, \quad v_1 - v_2 = 0, \quad s_1 - s_2 = 0,$$

so $P_{\mathcal{K}} \subset (\mathbb{P}^5)^2$ is cut out by five independent linear constraints.

Under $\pi_{\omega^*}$ we have $K_1 = K_2 \Rightarrow \omega_1^* \sim \omega_2^*$. With the scale of each DIAC fixed (e.g. $(\omega_i^*)_{33} = 1$ for $i = 1, 2$), this becomes the five linear constraints

$$o_{11} - o_{21} = 0,\ o_{12} - o_{22} = 0,\ o_{13} - o_{23} = 0,\ o_{14} - o_{24} = 0,\ o_{15} - o_{25} = 0.$$

Hence

$$P_{\mathcal{W}} = \{(\omega_1^*, \omega_2^*) \in (\mathbb{P}^5)^2 \mid \omega_1^* \sim \omega_2^*\} \cong \mathbb{P}^5 \quad \Rightarrow \quad \dim(P_{\mathcal{W}}) = 5.$$

In particular, although $(\mathbb{P}^5)^2$ has dimension 10, the constant-camera prior reduces the DIAC degrees of freedom to five.

### 4.2 Complexity of Minimal Self-Calibrations with Varying Intrinsics

Our goal is to characterise the algebraic complexity of varying-intrinsic Kruppa calibration tasks by analysing how many solutions their associated polynomial systems admit. Intuitively, tasks that produce many complex solutions are intrinsically more ambiguous and therefore more sensitive to noise and initialisation. To quantify this effect, we compute algebraic bounds on the number of solutions of generic minimal subsystems.

*Calibration Priors and Constant Assumptions across views.* In Tab. 1, a *task* specifies linear priors on intrinsics (e.g. $s_i = 0$, $f_i = g_i$, or other parameter eliminations) and certain parameters are *shared* across views, thereby defining a feasible intrinsic set $P_{\mathcal{K}}$ and the induced DIAC-feasible set $P_{\mathcal{W}}$. For example, `Task=ff000` with `Shared = f` (which we call "`ff000-f`" later) enforces $f_1 = f_2$ between two views, whereas `Shared=none` allows focal length vary independently. Each pair (Task, Shared) thus uniquely determines a pushforward variety $P_{\mathcal{W}}$.

*Sampling Synthetic scenes.* For each task and view count $m$, we generate *generic* problem instances by sampling $(K_i, R_i, T_i)$ consistent with the task prior (*e.g.* zero skew $s = 0$.) From these problem instances, we compute fundamental matrices and epipoles which are exact (up to machine precision), e.g. with the formulas

$$F_{ij} \ \sim \ K_j^{-\top} R_j^{-\top} [T_j - T_i]_\times R_i^{-1} K_i^{-1}. \tag{6}$$

Thus, $(F_{ij}, e'_{ij})$ are fixed in Kruppa's equations. Only $\omega^*$ is unknown.

*0-dimensional subsystems for root counting.* Following Sec. 3.3, we construct square 0-dimensional systems by evaluating the full set of constraints (Kruppa and pushforward) on generic fabricated scenes and extracting locally independent subsets via Jacobian rank at the ground-truth DIACs. These locally minimal relaxations are regular at the sampled generic instance (full Jacobian rank) and therefore define locally 0-dimensional systems with isolated solutions. We sample $(N = 10)$ minimal relaxations for each (Task, Shared) configuration and estimate the algebraic root counts reported below.

*Measuring Complexity of Self-calibration Systems.* For the ten square subsystems defining minimal relaxations for each task, shared-variable configuration, and minimal view count, we report the minimum and maximum values for one of three measures of algebraic complexity:

- **Root count:** the number of isolated complex solutions of a minimal relaxation, estimated via numerical monodromy heuristics [4, 7].
- **BKK bound:** the Bernstein–Kushnirenko–Khovanskii bound [25, Ch. 8.5], based on the mixed volume of a square subsystem's Newton polytopes.
- **Bézout bound:** the product of equations' degrees in a square subsystem.

These quantities satisfy the following inequalities [25, Ch. 8]:

$$\text{Root count} \leq \text{BKK} \leq \text{Bézout}.$$

However, the cost of computing each is often in the reverse order. Accordingly, we have reported exact root counts whenever this was possible with our limited computational resources. In cases where monodromy did not stabilize after finding 100K solutions, we have instead reported the BKK bound as a proxy for algebraic complexity. Similarly, in challenging cases where computing a large BKK bounds triggered an overflow errors, we fall back to the Bézout bound.

*Results* Table 1 summarises the empirical root counts and polyhedral (BKK) bounds for representative (Task, Shared) configurations. A clear trend emerges: stronger intrinsic coupling across views (i.e. more shared variables) often substantially reduces both the observed number of isolated complex solutions and the corresponding BKK bounds. In contrast, weakly constrained varying-intrinsic tasks typically exhibit rapid growth in their generic root counts (or, when applicable, upper bounds thereof.) This confirms that the algebraic complexity of self-calibration depends critically on the intrinsic prior, not on view count alone.

**Table 1:** Algebraic complexity estimates for varying-intrinsic self-calibration. Cells are colour-coded by the source of the reported solution count/bound. RootCount rows indicate monodromy estimates for #sol; other rows indicate fallbacks BKK and Bézout.

| Task | Shared-Vars | #views | dim($P_{\mathcal{W}}$) | #Subsys | Max #sol. | Min #sol. |
|---|---|---|---|---|---|---|
| ff000 | f | 2 | 1 | 1.5e1 | 2 | 1 |
| ff000 | none | 2 | 2 | 1.05e2 | 1 | 1 |
| ff00s | fs | 2 | 2 | 1.05e2 | 2 | 1 |
| ff00s | none | 3 | 6 | 8.15e6 | 6 | 1 |
| ffuv0 | fuv | 3 | 3 | 1.42e4 | 16 | 16 |
| ffuv0 | none | 5 | 15 | 1.62e20 | 372918 | 372191 |
| ffuv0 | uv | 3 | 5 | 1.22e6 | 40 | 40 |
| ffuvs | fuv | 3 | 6 | 8.15e6 | 24794911296 | 24794911296 |
| ffuvs | fuvs | 3 | 4 | 1.49e5 | 27 | 25 |
| ffuvs | none | 3 | 12 | 2.88e10 | 262144 | 262144 |
| ffuvs | uvs | 3 | 6 | 8.15e6 | 330225942528 | 330225942528 |
| fg000 | fg | 2 | 2 | 1.05e2 | 3 | 3 |
| fg000 | none | 3 | 6 | 8.15e6 | 3 | 1 |
| fg00s | fg | 3 | 5 | 1.22e6 | 5 | 1 |
| fg00s | none | 5 | 15 | 1.62e20 | 56 | 35 |
| fguv0 | fg | 4 | 10 | 5.72e12 | 61917364224 | 61917364224 |
| fguv0 | fguv | 3 | 4 | 1.49e5 | 14 | 14 |
| fguv0 | uv | 4 | 10 | 5.72e12 | 144 | 98 |
| fguvs | fg | 6 | 20 | 1.9e28 | 8153726976 | 8153726976 |
| fguvs | fguv | 4 | 8 | 7.75e10 | 4294967296 | 4294967296 |
| fguvs | fguvs | 3 | 5 | 1.22e6 | 18 | 17 |
| fguvs | uv | 6 | 20 | 1.9e28 | 6060 | 4411 |
| fguvs | uvs | 4 | 11 | 4.16e13 | 82644187136 | 82644187136 |

These findings highlight the limitations of classical degrees-of-freedom (DoF) counting arguments [11, 16] for determining "minimal views." Such analyses assume that each additional constraint reduces the solution-set dimension by one, i.e. that the system behaves as a complete intersection with generically independent constraints. For Kruppa-based formulations with varying intrinsics, this assumption is fragile: the projective nature of Kruppa relations and task-induced structure in DIAC space can introduce algebraic dependencies, which may invalidate naive dimension-counting and/or upper bounds on generic root counts.

Moreover, minimal view counts are not directly comparable across formulations. For instance, [24] performs metric upgrade via the dual absolute quadric (DAQ), introducing additional unknowns beyond the per-view DIACs considered here which must be estimated using more than a minimal number of views.

## 5 Experiments

This section evaluates the practical impact of the proposed minimal relaxations. We assess their numerical stability and calibration accuracy on synthetic and real-world data, demonstrating that the selected locally independent subsystems remain robust under noise and varying-intrinsic configurations. Our solvers are intended as algebraic initializers and analysis tools: in a full SfM pipeline, their estimates would typically be followed by non-linear refinement of camera poses, intrinsics, and structure. Implementation details and bundle-adjustment results are provided in SM. D.2 and SM. D.8, respectively.

*Tasks.* We evaluate the following tasks (with the minimal number of views used in our solver instantiations):

– `ff000-f` (2 views): only focal length unknown and *shared* across views.
– `ff000-none` (2 views): only focal length unknown, not shared.
– `ffuv0-uv` (3 views): skew 0, fixed aspect ratio; principal point shared $(u, v)$.
– `fguv0-uv` (4 views): skew 0; principal point shared $(u, v)$; focal lengths vary.

*Solvers.* (i) **Kruppa (Ours)**: For each task (with $N = 10$ sampled minimal relaxations), we solve for $\omega^* = (\omega_1^*, \dots, \omega_m^*)$ using the Kruppa relations (1) together with the pushforward DIAC constraints defining $P_{\mathcal{W}}$ for the given task. We form a square system via precomputed minimal relaxations (Sec. 3.3) and recover $K_i$ from $\omega_i^* \sim K_i K_i^\top$ (Cholesky/RQ factorisation with the task constraints enforced by construction). From all solutions to sampled relaxations, we choose the solution whose $K$ matrix is closest to the ground-truth.

(ii) **Linear DAQ Solver [24]**: For `ff000-f`, we implement the linear DAQ-based initialisation of Pollefeys et al. [24], which was designed specifically for the shared-focal setting. We use the same noisy $F_{ij}$ estimated from correspondences and recover the intrinsics via their linear constraints on the dual absolute quadric, followed by metric upgrade. This baseline is not directly applicable to the other varying-intrinsic tasks considered here, so we report it only for `ff000-f`.

(iii) **Pairwise Initialisation (Lourakis [22])**: For each image pair $(i, j)$, two Kruppa-derived polynomial equations in $(f_i^2, f_j^2)$ are reduced to a homogeneous linear system and solved via SVD. Physically invalid solutions (non-real or non-positive) are discarded, and the focal length of each view is initialised as the median of all valid pairwise estimates involving that view. Principal points are initialised at the image centre (as COLMAP's guess in Sec. 5.2). This procedure provides a robust per-view initialisation but does not enforce global consistency across multiple views. Thus, we report it for `ff000-none` and `ffuv0-uv`.

*Metrics.* We report relative errors averaged over views and then averaged over sampled random scenes at each noise level:

$$\Delta f = \frac{|\hat{f} - f|}{f}, \quad \Delta g = \frac{|\hat{g} - g|}{g}, \quad \Delta u = \frac{|\hat{u} - u|}{\max(1, |u|)}, \quad \Delta v = \frac{|\hat{v} - v|}{\max(1, |v|)}.$$

For compact comparison, we summarise focal-length and principal-point performance by $\Delta fg := \frac{1}{2}(\Delta f + \Delta g)$ and $\Delta uv := \frac{1}{2}(\Delta u + \Delta v)$. (For tasks where $(u, v)$ are fixed by the prior, $\Delta uv$ is effectively zero up to numerical precision.)

### 5.1 Synthetic Images: Varying-Intrinsic Self-Calibration Tasks

We first evaluate calibration accuracy on synthetic image correspondences under controlled pixel noise, comparing multiple intrinsic-task configurations. A key reference point is the classical case `ff000` with a shared focal length across two views, where our Kruppa–DIAC formulation is expected to be comparable to the linear DAQ initialisation of Pollefeys et al [24]. Results with additional solvers can be found in SM. D.5.

*Setup: Scene generation.* For each noise level $\sigma \in \{0.0, 0.1, \ldots, 1.0\}$, we generate 100 random scenes. Each scene is defined by sampling intrinsics $K_i$, rotations $R_i$ (Cayley parameterisation), and translations $T_i$ consistent with the task prior (shared variables, zero skew, fixed aspect, etc.). We sample 1000 random 3D points and project them into each view using

$$P_i = K_i R_i [I \mid -T_i].$$

We corrupt 2D image measurements with i.i.d. Gaussian pixel noise of standard deviation $\sigma$ and estimate the fundamental matrices $F_{ij}$ (and epipoles $e'_{ij} \in \ker(F_{ij}^\top)$) using the normalised 8-point algorithm on noisy correspondences [13], which provides improved numerical stability compared to the 7-point variant [5]. 2D image measurements are scaled appropriately SM. D.1.

**Results.** Figure 1 compares $\Delta fg$ and $\Delta uv$ across noise levels.

*Ours vs. Baselines.* On the shared-focal two-view task `ff000-f`, our Kruppa–DIAC solver is on par with the linear DAQ baseline: both exhibit comparable $\Delta fg$ across $\sigma \in [0.1, 1.0]$ (e.g. at $\sigma = 1.0$, $\Delta fg = 0.511$ (ours) vs. 0.457 (DAQ)). For `ff000-none` and `ffuv0-uv`, the pairwise scheme of Lourakis is noticeably less accurate under noise in our setting. This is unsurprising since it was proposed as an initialisation for subsequent non-linear refinement, whereas we directly solve globally coupled multiview DIAC constraints.

*Task difficulty under varying intrinsics.* Moving beyond the classical `ff000` regime, task difficulty diverges substantially. The fixed-aspect, shared-principal-point task `ffuv0-uv` is markedly harder, with errors increasing rapidly with noise (e.g. at $\sigma = 1.0$, $\Delta fg = 1.737$ and $\Delta uv = 1.507$). In contrast, the more realistic `fguv0-uv` setting (varying focal lengths, shared principal point, zero skew) remains stable across noise with moderate degradation (e.g. at $\sigma = 1.0$, $\Delta fg = 0.547$ and $\Delta uv = 0.413$). These trends align with our algebraic complexity analysis: weaker intrinsic coupling leads to both greater solution ambiguity and greater sensitivity to noisy $F_{ij}$ (Tab. 1).

### 5.2 Real-World Images: Varying Focal Length

We run a real-world evaluation designed to test robustness under genuinely dynamic intrinsics. We use the InFlux dataset [21], which provides real video

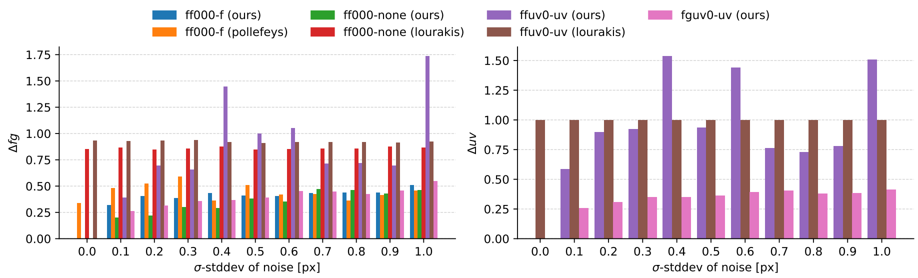


**Fig. 1:** Synthetic image experiments. Left: focal-length error $\Delta fg$ versus pixel noise. Right: principal-point error $\Delta uv$. We compare our solvers across tasks to the baselines.

sequences with frame-level ground-truth intrinsic parameters recorded during capture (SM. D.6): *bike_ shot6*, *bike_ shot11*, *friend_ 1_ shot3*, *statue_ shot2*, and *statue_ shot6*. These sequences exhibit observable focal variation and moderate principal-point drift. Sequences with negligible intrinsic change are excluded, as they do not meaningfully stress the varying-intrinsic regime. Table 2 reports results averaged over these five sequences; the full per-sequence breakdown is provided in SM. D.7.

*Protocol and robust estimation.* For each video, we sample subsets of $m = 10$ sharp frames exhibiting focal variation and extract 2D correspondences for all candidate view pairs. Fundamental matrices $F_{ij}$ are estimated using the normalised 8-point algorithm as in the synthetic experiments (Sec. 5.1). All solvers are evaluated within an MSAC framework [28]: at each iteration, a random 8-point subset is sampled, the corresponding $F_{ij}$ are estimated, and the solution yielding the lowest sum of intrinsic errors is retained over 20 iterations. No sequence-specific tuning is performed; the solver configuration and relaxation bank are identical to the synthetic experiments.

*Results.* Compared to the baselines, two observations emerge. (i) `ff000-f` (static-focal setting), our Kruppa solver performs comparably to the linear DAQ baseline, confirming consistency with classical shared-focal self-calibration in real data, (ii) `ffuv0-uv`, while the pairwise initialisation method of Lourakis achieves comparable focal-length accuracy, its principal-point estimates are substantially less stable in our setting due to its image centre assumption.

Under these genuinely varying-intrinsic regimes, our formulation remains stable despite real correspondence noise and imperfect epipolar estimation. Consistent with the synthetic study, tasks imposing stronger intrinsic coupling yield more reliable principal-point estimates, while focal parameters remain bounded even under substantial zoom variation. Bundle-adjustment results using COLMAP are reported in SM. D.8.

### 5.3 Real-World Images: Edited frames

Many consumer video pipelines apply electronic image stabilisation (EIS) by dynamically cropping and shifting the image window to compensate for camera

**Table 2:** Sequence-level aggregation (unweighted across sequences). For each sequence, we take the median error over frames, then report mean±std across sequences. Lower is better.

| Method | $\Delta fg$ (mean±std) ↓ | $\Delta uv$ (mean±std) ↓ |
|---|---|---|
| `ff000-f (Kruppa)` | $0.725 \pm 0.169$ | – |
| `ff000-f (Pollefeys)` | $0.857 \pm 0.116$ | – |
| `ff000-none` | $0.604 \pm 0.162$ | – |
| `ffuv0-uv (Kruppa)` | $0.790 \pm 0.104$ | $0.312 \pm 0.086$ |
| `ffuv0-uv (Lourakis)` | $0.805 \pm 0.140$ | $0.797 \pm 0.001$ |
| `fguv0-uv` | $0.748 \pm 0.105$ | $0.464 \pm 0.133$ |

shake, which alters the pixel coordinate system of the exported frames and effectively shifts the principal point away from the image centre. Consequently, centre-based initialisation heuristics—commonly used in practice—become systematically biased. We evaluate robustness to this realistic coordinate perturbation.

*Setup.* Starting from the same five InFlux sequences used in Sec. 5.1, we simulate a realistic electronic image stabilisation (EIS) effect by cropping each frame to a fixed $1400 \times 800$ window anchored at the bottom-right of the original image. This preserves the underlying optics but shifts the principal point in pixel coordinates relative to the new image centre. Ground-truth intrinsics are updated accordingly by subtracting the crop offset from $(u, v)$. All other components (correspondences, normalised 8-point $F_{ij}$ estimation with MSAC, solver settings, and precomputed relaxations) follow Sec. 5.2 unchanged.

*Baseline.* We compare against a centre-based initialisation commonly used in practice: *COLMAP-guess*, which sets $f = 1.2 \max(W, H)$ and $(u, v) = \left(\frac{W}{2}, \frac{H}{2}\right)$ where $(W, H)$ denote the cropped image dimensions. This heuristic assumes the principal point lies at the image centre. The pairwise initialisation of Lourakis [22] also relies on the same centre-principal-point assumption; therefore we do not include it in this experiment.

*Results.* Table 3 shows that centre-based initialisation fails dramatically in principal-point estimation, with $\Delta uv = 2.483 \pm 0.608$. This error magnitude reflects the systematic bias introduced by the crop offset rather than random noise. In contrast, both Kruppa-based tasks recover substantially more accurate principal points ($\Delta uv = 0.395 \pm 0.475$ for `ffuv0-uv` and $0.665 \pm 0.106$ for `fguv0-uv`), demonstrating that solving the algebraic DIAC constraints is robust to mis-centred images. Focal-length accuracy remains comparable across methods in this edited setting, since cropping alters the image origin but does not change the focal length in pixel units. Consequently, centre-based heuristics mainly fail in estimating $(u, v)$ rather than $(f, g)$.

**Table 3:** Sequence-level aggregation. Mean per sequence, then mean±std across sequences.

| Task | $\Delta fg$ (mean±std) ↓ | $\Delta uv$ (mean±std) ↓ |
|---|---|---|
| `COLMAP-guess` | $0.692 \pm 0.126$ | $2.483 \pm 0.608$ |
| `ffuv0-uv` | $0.857 \pm 0.108$ | $0.395 \pm 0.475$ |
| `fguv0-uv` | $0.890 \pm 0.033$ | $0.665 \pm 0.106$ |

## 6 Conclusion

Our formulation of varying-intrinsic self-calibration, combining pushforward DIAC constraints and minimal relaxation, provides a unified approach to calibration with partial priors, including several less-studied cases. Our experiments indicate that structured intrinsic coupling can reduce algebraic ambiguity and improve robustness in the tested generic settings. The resulting solvers are best viewed as algebraic initialisers and analysis tools; a complete treatment of critical motions, artificial degeneracies, and specialised solver comparisons remains future work.

## References

1. Barath, D., Toth, T., Hajder, L.: 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. pp. 6003–6011 (2017) 3
2. Bocquillon, B., Bartoli, A., Gurdjos, P., Crouzil, A.: On constant focal length self-calibration from multiple views. In: 2007 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1–8. IEEE (2007) 3
3. Bougnoux, S.: From projective to euclidean space under any practical situation, a criticism of self-calibration. In: Sixth International Conference on Computer Vision (IEEE Cat. No. 98CH36271). pp. 790–796. IEEE (1998) 3
4. Breiding, P., Timme, S.: Homotopycontinuation. jl: A package for homotopy continuation in julia. In: International congress on mathematical software. pp. 458–465. Springer (2018) 9
5. Cin, A.P.D., Duff, T., Magri, L., Pajdla, T.: Minimal perspective autocalibration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5064–5073 (2024) 2, 5, 6, 12
6. Cox, D., Little, J., O'shea, D., Sweedler, M.: Ideals, varieties, and algorithms. Springer (1997) 8
7. Duff, T., Hill, C., Jensen, A., Lee, K., Leykin, A., Sommars, J.: Solving polynomial systems via homotopy continuation and monodromy. IMA Journal of Numerical Analysis **39**(3), 1421–1446 (2019) 9
8. Duff, T., Kohn, K., Leykin, A., Pajdla, T.: PLMP-point-line minimal problems in complete multi-view visibility. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1675–1684 (2019) 5
9. Fetzer, T., Reis, G., Stricker, D.: Stable intrinsic auto-calibration from fundamental matrices of devices with uncorrelated camera parameters. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 221–230 (2020) 3

10. Grayson, D.R., Stillman, M.E.: Macaulay2, a software system for research in algebraic geometry. Available at http://www2.macaulay2.com 8
11. Hartley, R.I., Zisserman, A.: Multiple View Geometry in Computer Vision. Cambridge University Press, ISBN: 0521540518, second edn. (2004) 1, 3, 4, 10
12. Hartley, R.I.: Estimation of relative camera positions for uncalibrated cameras. In: European conference on computer vision. pp. 579–587. Springer (1992) 3
13. Hartley, R.I.: In defense of the eight-point algorithm. IEEE Transactions on pattern analysis and machine intelligence **19**(6), 580–593 (1997) 12
14. Hartley, R.I.: Kruppa's equations derived from the fundamental matrix. IEEE Transactions on pattern analysis and machine intelligence **19**(2), 133–135 (1997) 2, 4, 8
15. Heikkila, J.: Using sparse elimination for solving minimal problems in computer vision. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 76–84 (2017) 3
16. Hemayed, E.E.: A survey of camera self-calibration. In: Proceedings of the IEEE Conference on Advanced Video and Signal Based Surveillance, 2003. pp. 351–357. IEEE (2003) 1, 2, 10
17. Heyden, A., Åkström, K.: Minimal conditions on intrinsic parameters for euclidean reconstruction. In: Asian Conference on Computer Vision. pp. 169–176. Springer (1998) 1, 2
18. Heyden, A., Astrom, K.: Euclidean reconstruction from image sequences with varying and unknown focal length and principal point. In: Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 438–443. IEEE (1997) 1, 2
19. Huang, T.S., Faugeras, O.D.: Some properties of the e matrix in two-view motion estimation. IEEE Transactions on Pattern Analysis and Machine Intelligence **11**(12), 1310–1312 (2002) 3
20. Kosaka, N., Duff, T., Pajdla, T.: Minimal constraint relaxation for multiview autocalibration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 28937–28946 (June 2026) 2
21. Liang, E., Bhattacharjee, R., Dey, S., Moschopoulos, R., Wang, C., Liao, M., Tan, G., Wang, A., Kayan, K., Alexandropoulos, S., et al.: Influx: A benchmark for self-calibration of dynamic intrinsics of video cameras. arXiv preprint arXiv:2510.23589 (2025) 12
22. Lourakis, M.I., Deriche, R.: Camera self-calibration using the Kruppa equations and the SVD of the fundamental matrix: The case of varying intrinsic parameters. Ph.D. thesis, INRIA (2000) 2, 3, 4, 11, 14
23. Luong, Q.T., Faugeras, O.D.: Self-calibration of a moving camera from point correspondences and fundamental matrices. International Journal of computer vision **22**(3), 261–289 (1997) 2, 4, 8
24. Pollefeys, M., Koch, R., Gool, L.V.: Self-calibration and metric reconstruction inspite of varying and unknown intrinsic camera parameters. International journal of computer vision **32**(1), 7–25 (1999) 1, 2, 3, 10, 11, 12
25. Sommese, A.J., Wampler, C.W., et al.: The Numerical solution of systems of polynomials arising in engineering and science. World Scientific (2005) 6, 9
26. Sturm, P.: Critical motion sequences for monocular self-calibration and uncalibrated euclidean reconstruction. In: Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 1100–1105. IEEE (1997) 3, 6

27. Sturm, P.: On focal length calibration from two views. In: IEEE International Conference on Computer Vision and Pattern Recognition (CVPR'01). vol. 2, pp. 145–150. IEEE Computer Society (2001) 3
28. Torr, P.H., Zisserman, A.: Mlesac: A new robust estimator with application to estimating image geometry. Computer vision and image understanding **78**(1), 138–156 (2000) 13