

Galois/monodromy groups for decomposing minimal problems in 3D reconstruction

Timothy Duff,^{*}
Viktor Korotynskiy,[†]
Tomas Pajdla,[†]
Margaret H. Regan[‡]

Abstract

We consider Galois/monodromy groups arising in computer vision applications, with a view towards building more efficient polynomial solvers. The Galois/monodromy group allows us to decide when a given problem decomposes into algebraic subproblems, and whether or not it has any symmetries. Tools from numerical algebraic geometry and computational group theory allow us to apply this framework to classical and novel reconstruction problems. We consider three classical cases—3-point absolute pose, 5-point relative pose, and 4-point homography estimation for calibrated cameras—where the decomposition and symmetries may be naturally understood in terms of the Galois/monodromy group. We then show how our framework can be applied to novel problems from absolute and relative pose estimation. For instance, we discover new symmetries for absolute pose problems involving mixtures of point and line features. We also describe a problem of estimating a pair of calibrated homographies between three images. For this problem of degree 64, we can reduce the degree to 16, the latter better reflecting the intrinsic difficulty of algebraically solving the problem. As a byproduct, we obtain new constraints on compatible homographies, which may be of independent interest.

Keywords. Algebraic vision, 3D reconstruction, minimal problems, Galois groups, monodromy groups, numerical algebraic geometry

AMS Subject Classification: 12F10, 14Q15, 65H14, 68T45

1 Introduction

Estimating 3D geometry from data [38, 84, 87] in one or more images is a reoccurring problem in subjects like computer vision, robotics, and photogrammetry. The exact formulations of these problems vary considerably, depending on factors like the model of image formation that is being considered or what data are available. A common element in many of these geometric pose estimation problems is the presence of polynomial constraints, depending on both the data and the unknown quantities to be estimated. Considerable emphasis has been placed on *minimal problems* which are usually well-posed in the sense that exact solutions exist for generic data. A thorough understanding of these problems is not only theoretically appealing, but has practical consequences. Minimal

^{*}TD: University of Washington. timduff@uw.edu

[†]VK, TP: Czech Institute of Informatics, Robotics and Cybernetics, Czech Technical University in Prague, CZ. viktor.korotynskiy@cvut.cz, pajdla@cvut.cz

[‡]MHR: Duke University. mregan@math.duke.edu

problems and their solvers [1, 8, 9, 10, 15, 16, 20, 23, 24, 26, 37, 48, 50, 52, 53, 54, 56, 57, 58, 67, 69, 72, 77, 82, 83, 90, 93] play an outsized role in RANSAC-based estimation, introduced to the computer vision community by Fischler and Bolles in 1981 [28] and later developed to a very efficient and robust estimation method [76]. Using a robust minimal solver as a subroutine within the RANSAC loop allows for fewer iterations when compared to other techniques (eg. linear algorithms [39]) that require more measurements and often ignore the underlying polynomial constraints. However, the algebraic complexity of a minimal problem as measured by its *degree* has been observed as a limiting factor for minimal solvers, despite their many successes in practice.

Towards controlling this algebraic complexity, the question of how to detect and exploit symmetries when solving polynomial systems of equations has attracted interest in recent literature in computer vision [6, 55]. The techniques developed there can be understood in the context of linear representation theory, and can be traced back to pioneering work in computer algebra [17, 31]. As noted in [55, Sec 4], these techniques may be difficult to apply without some problem-specific knowledge. Moreover, many problems have *nonlinear* symmetries. A famous example is the problem of 5-point relative pose estimation, or simply the 5-point problem. The set of solutions to this problem is invariant under a nonlinear symmetry known as the twisted pair.

Example 1.1. Two cameras view some number of points in the world (ie. 3-dimensional space.) Each camera is modeled via perspective projection. When the cameras are calibrated [38, Ch. 6], we may assume that the two camera frames differ by a rotation $\mathbf{R}$ and a translation $\mathbf{t}$.




Figure 1: Schematic depiction of the five-point problem

In the five-point problem, we are given the 2D data of 5 correspondences between image points $\mathbf{x}_1 \leftrightarrow \mathbf{y}_1, \dots, \mathbf{x}_5 \leftrightarrow \mathbf{y}_5$ which are assumed to be images of 5 world points under two calibrated cameras. The image points are given in *normalized image coordinates*, meaning that each is a 3×1 vector whose last coordinate equals 1. The task is to reconstruct the relative orientation $[\mathbf{R} \mid \mathbf{t}] \in \text{SE}_{\mathbb{R}}(3)$ between the two views and each of the five world points as measured by their

depths with respect to the first and second camera frames. Writing $\alpha_1, \dots, \alpha_5$ for the depths with respect to the first camera and $\beta_1, \dots, \beta_5$ for the depths with respect to the second camera, solutions to the five-point problem must satisfy a system of polynomial equations and inequations:

$$\begin{aligned} \mathbf{R}^\top \mathbf{R} &= \mathbf{I}, \quad \det \mathbf{R} = 1, \\ \beta_i \mathbf{y}_i &= \mathbf{R} \alpha_i \mathbf{x}_i + \mathbf{t}, \quad \alpha_i, \beta_i \neq 0, \quad \forall i = 1, \dots, 5. \end{aligned} \tag{1}$$

An inherent ambiguity of the five-point problem is that the unknowns $\mathbf{t}, \alpha_1, \dots, \alpha_5, \beta_1, \dots, \beta_5$ can only be recovered up to a common scale factor. If we treat these unknowns as homogeneous coordinates on a 12-dimensional projective space, then for generic data $\mathbf{x}_1, \dots, \mathbf{x}_5, \mathbf{y}_1, \dots, \mathbf{y}_5$, there are at most finitely many solutions in $(\mathbf{R}, \mathbf{t}, \alpha_1, \dots, \alpha_5, \beta_1, \dots, \beta_5) \in \text{SO}_{\mathbb{C}}(3) \times \mathbb{P}_{\mathbb{C}}^{12}$ to the system (1). In practice, these solutions may be computed by working in a fixed affine patch of $\mathbb{P} = \mathbb{P}_{\mathbb{C}}^{12}$ (eg. $\alpha_1 = 1$.)

If we count solutions over the complex numbers, there are exactly 20 solutions for generic data in $Z = (\mathbb{C}^2 \times \{1\})^5 \times (\mathbb{C}^2 \times \{1\})^5$. The solutions to (1) are naturally identified with the fibers of a branched cover $f : X \rightarrow Z$ (see Definition 2.1), where X is the *incidence correspondence*

$$X = \{(\mathbf{R}, (\mathbf{t}, \alpha_1, \dots, \alpha_5, \beta_1, \dots, \beta_5), (\mathbf{x}_1, \dots, \mathbf{x}_5, \mathbf{y}_1, \dots, \mathbf{y}_5)) \in \text{SO}_{\mathbb{C}}(3) \times \mathbb{P}_{\mathbb{C}}^{12} \times Z \mid (1) \text{ holds} \}.$$

Fixing generic data $(\mathbf{x}_1, \dots, \mathbf{y}_5) \in Z$, we may associate to each solution to (1) a *twisted pair* solution, obtained by rotation of the second camera frame by 180° about the baseline connecting the first and second camera centers.




Figure 2: Twisted pair symmetry of the five-point problem

The twisted pair may be viewed as a rational map $\Psi: X \dashrightarrow X$, given coordinate-wise by

$$\begin{aligned}\Psi(\mathbf{R}) &= \left(\frac{2}{\|\mathbf{t}\|^2} \mathbf{t} \mathbf{t}^\top - \mathbf{I} \right) \mathbf{R} \\ \Psi(\mathbf{t}) &= \mathbf{t} \\ \Psi(\alpha_i) &= \frac{-\alpha_i \|\mathbf{t}\|^2}{\|\mathbf{t}\|^2 + 2 \langle \mathbf{R}^\top \mathbf{t}, \alpha_i \mathbf{x}_i \rangle} = \frac{-\alpha_i \|\mathbf{t}\|^2}{\|\beta_i \mathbf{y}_i\|^2 - \|\alpha_i \mathbf{x}_i\|^2} \\ \Psi(\beta_i) &= \frac{\beta_i \|\mathbf{t}\|^2}{\|\mathbf{t}\|^2 + 2 \langle \mathbf{R}^\top \mathbf{t}, \beta_i \mathbf{x}_i \rangle} = \frac{\beta_i \|\mathbf{t}\|^2}{\|\beta_i \mathbf{y}_i\|^2 - \|\alpha_i \mathbf{x}_i\|^2} \\ \Psi(\mathbf{x}_1, \dots, \mathbf{x}_5, \mathbf{y}_1, \dots, \mathbf{y}_5) &= (\mathbf{x}_1, \dots, \mathbf{x}_5, \mathbf{y}_1, \dots, \mathbf{y}_5).\end{aligned}\tag{2}$$

Here, we use the notation $\langle, \rangle$ and $\|\cdot\|^2$ for the complex quadratic forms $\langle \mathbf{a}, \mathbf{b} \rangle = \mathbf{a}^\top \mathbf{b}$, $\|\mathbf{a}\|^2 = \langle \mathbf{a}, \mathbf{a} \rangle$ which restrict to the usual norm and inner product on $\mathbb{R}^3$. We note that Ψ is undefined whenever $\mathbf{t} \in \mathbb{P}^2$ is an *isotropic vector* satisfying $\|\mathbf{t}\|^2 = 0$, and whenever $\|\alpha_i \mathbf{x}_i\|^2 = \|\beta_i \mathbf{y}_i\|^2$ for some $i = 1, \dots, 5$. The second condition can be understood geometrically: if the camera centers and the world point $\mathbf{X} = \alpha \mathbf{x}$ form an isosceles triangle with base $\|\mathbf{t}\|$, then, after rotating the second camera, the rays which join camera centers to the respective image points will become parallel. When the map Ψ is defined, these two rays meet in a new world point $\mathbf{X}'$, as depicted in Figure 2.

One can check (e.g. [65, p. 20]) that the set of solutions to equations (1) is left invariant under application of Ψ . In other words, we have an equality of mappings $f \circ \Psi = f$ wherever Ψ is defined. The map Ψ is a deck transformation (see Definition 2.12) of the branched cover f .

Readers familiar with the *four-fold symmetry* depicted in Hartley and Zisserman’s classic text (see [38, Figure 9.12]) may wonder how this relates to the branched cover described above. We offer an explanation in Section 4.1.

The twisted pair of Example (1) is a well-known construction. It is not easy to determine that such a symmetry exists just from staring at Equations (1). However, the existence of such a symmetry can be easily decided after computing the Galois/monodromy group of f , as we make explicit in Proposition 2.16. The state-of-the-art method [72] for solving the five-point problem is based on *decomposing* the branched cover f in terms of the *essential matrix* (see Examples 2.9, 2.15 and Section 4.1). In general, decomposability can be detected after computing the Galois/monodromy group, by Proposition 2.10.

Working within the framework of branched covers and Galois/monodromy groups allows us to probe the structure of minimal problems in ways that the previously mentioned works cannot. As is well-known (see Proposition 2.3), this group can be understood both algebraically (“Galois”) and topologically (“monodromy”). The algebraic point of view was pursued by Hartley, Nistér, and Stewénus [71], who computed Galois groups symbolically to show that certain formulations of minimal problems were degree-optimal. In our paper, the topological point of view is somewhat more relevant, since we can compute the monodromy action for problems of potentially large degree using numerical homotopy continuation methods, as implemented in several software packages [11, 13, 22].

Computing the Galois/monodromy group allows us to decide decomposability or the existence of deck transformations by reduction to standard algorithms in computational group theory [44]. These algorithms are implemented in computer algebra systems such as GAP [30], which was an essential ingredient in discovering our main results. We find frequently that many minimal problems do not decompose—in this case, the Galois/monodromy group tells us that our formulation is optimal in a precise sense. However, we find in several cases that minimal problems previously considered in the literature do decompose. We describe this in the context of problems where this phenomenon

is already well-understood (such as the five-point problem), and also in several cases where some decomposition or symmetry was not previously noticed.

There is a long history in computer vision of exploiting symmetries and decompositions of algebraic problems. For example, the classical approach to the five-point problem decomposes the 3D reconstruction problem of degree 20 into a sequence of two problems of degree two and degree 10. This brings about two advantages. First, the solutions are found from the eigenvectors of a 10×10 matrix, instead of 20×20 matrix. Non-real solutions from this first stage are regarded as spurious and may be discarded. The remaining solutions extend to solutions of the degree 20 problem after solving simple quadratic equations. Although this decomposition of the five-point problem is well-known, the procedure used to obtain it is *ad-hoc*, relying on some geometric insight into the problem. The previous-mentioned works of Ask et al. and Larsson et al. [6, 55] develop a systematic method for discovering a simple class of symmetries—see Example 2.14. These works also formulate a general approach for simplifying algebraic problems, analogous to that used for the five-point problem. In this approach, a large eigenvalue problem may be replaced by a sequence of smaller eigenvalue problems or, in some cases, by a single small eigenvalue problem and formulas that recover solutions of the original problem. Solutions of the smaller problems are screened as soon as possible, and we may discard those which do not satisfy additional problem-specific criteria (eg. semialgebraic constraints.) Galois/monodromy groups alone do not supply all of the essential ingredients needed to make this approach effective. However, they do precisely answer the question of when this approach can be used, and allow us to detect a much larger class of symmetries.

1.1 Main contributions and organization

Numerical techniques for computing Galois/monodromy groups are a previously-developed topic (see eg. [43, 61].) Our main contributions are based on applying these techniques to computer vision. Insight gained from these computations allow us to reveal structure within novel problems and explain structure that was previously observed in the literature. Our paper contains several original computational results: namely, Results 3.1, 4.1, 4.2, and 4.3. Since the methods used to establish these results may fall short of the conventional standard for proof, we use the term “Result” instead of “Theorem”. Code used to obtain these results, in addition to code illustrating many of the examples in the paper, may be found at the following url: <https://github.com/timduff35/galois-3d-vision>.

For readers with a background in computer vision, we suggest starting with the familiar examples in Sections 3 and 4, referring back to the theory and examples in Section 2 as needed. The mathematics needed to fully understand our paper is presented in Section 2.1. We include proofs of several standard facts in hope that our paper is accessible to a broad audience. We give a brief overview of numerical methods for computing Galois/monodromy groups in Section 2.2. In Section 3, we investigate the Galois/monodromy groups of absolute pose problems involving points and lines, where the task is to compute cameras from correspondence data between the world and an image. In particular, we describe newly discovered symmetries for problems involving a mixture of points and lines. Section 4 considers relative pose problems, where the task is to compute the transformation between camera frames given correspondence data between images. We describe decompositions for the five-point problem in Section 4.1, the four-point calibrated homography problem for two views in Section 4.2, and a novel minimal problem involving two calibrated homographies between three images in Section 4.3. We give a conclusion and outlook in Section 5.

2 Background

2.1 Branched covers and monodromy groups

Definition 2.1. A *branched cover* is a dominant, rational map $X \dashrightarrow Z$, where X and Z are irreducible algebraic varieties over $\mathbb{C}$ of the same dimension.

We emphasize several consequences of this definition. Most importantly, it follows that for generic (and hence *almost all*) data $z \in Z$, the fiber over z , denoted X_z , is a nonempty, finite set. Second is the assumption of irreducibility, which implies that the monodromy group acts transitively. In principle, we can always reduce to the case of an irreducible variety by considering the irreducible components of an arbitrary variety. The many examples we consider show that irreducibility is a natural assumption. Finally, although many of the branched covers we consider are actually regular maps, the maps appearing as deck transformations or in a factorization need not be defined on the same domain as the branched cover, making it more natural to work with rational maps. We say that X and Z are the *total space* and *base* of the branched cover, respectively. The reader may safely assume that all varieties are quasiprojective.

Pulling back rational functions from Z to X lets us identify $\mathbb{C}(Z)$ with a subfield of $\mathbb{C}(X)$. Since $\mathbb{C}(X)$ and $\mathbb{C}(Z)$ have the same transcendence degree over $\mathbb{C}$, the field extension $\mathbb{C}(X)/\mathbb{C}(Z)$ is finite. The *degree* of the map $X \dashrightarrow Z$ may be defined as the degree of this field extension. We write $\deg(X/Z)$ for this quantity, since the map $X \dashrightarrow Z$ is usually clear from context. We say that a nonempty Zariski-open $U \subset Z$ is a *regular locus* for $X \dashrightarrow Z$ if $U \cap Z_{\text{Sing}} = \emptyset$ and if for all $z \in U$ the cardinality of the fiber $X_z := f^{-1}(z)$ is equal to the degree of the map. The existence of such a U follows from basic results in algebraic geometry [86, cf. pp. 142].

We now recall the monodromy action on the fibers of a degree- d branched cover $f: X \dashrightarrow Z$. We omit many details which may be found in several excellent introductory references [40, 68, 95]. Fix a regular locus U and a basepoint $z \in U$, and write $X_z = \{x_1, \dots, x_d\}$. A *loop* based at z is a continuous map $\gamma: [0, 1] \rightarrow U$ which satisfies $\gamma(0) = \gamma(1) = z$. For each x_i , there exists a unique *lift* $\tilde{\gamma}_i: [0, 1] \rightarrow f^{-1}(U)$ satisfying $\gamma = f \circ \tilde{\gamma}_i$ and $\tilde{\gamma}_i(0) = x_i$. This fact from topology is known as the *unique path-lifting property* (see eg. [40, Proposition 1.30]). The lifts based at each of the points $x_1, \dots, x_d$ determine a permutation of the fiber, $\sigma_\gamma: X_z \rightarrow X_z$, which may be written in two-line notation as

$$\sigma_\gamma = \begin{pmatrix} x_1 & x_2 & x_3 & \cdots & x_d \\ \tilde{\gamma}_1(1) & \tilde{\gamma}_2(1) & \tilde{\gamma}_3(1) & \cdots & \tilde{\gamma}_d(1) \end{pmatrix}. \quad (3)$$

Remark 2.2. We use standard notation from group theory: $\text{Sym}(X)$ is the symmetric group of all permutations from a finite set X to itself, S_n , A_n , denote symmetric and alternating groups acting on letters $[n] = \{1, \dots, n\}$ (hence $\text{Sym}([n]) = S_n$), and C_n denotes a cyclic group of order n .

For completeness, we define the semidirect products and wreath products which figure prominently in Result 3.1 and Result 4.1. If G and H are two groups and

$$\begin{aligned} \varphi: H &\rightarrow \text{Aut}(G) \\ h &\mapsto \varphi_h \end{aligned}$$

is a homomorphism, then the *semidirect product* $G \rtimes_\varphi H$ is a group whose underlying set is the Cartesian product $G \times H$. Multiplication $\star$ in this group is defined by

$$(g_1, h_1) \star (g_2, h_2) = (g_1 \varphi_{h_1}(g_2), h_1 h_2).$$

The *wreath product* of $P \wr H$ of two permutation groups $P \subseteq S_m$ and $H \subseteq S_n$ is a special case of this construction. It may be defined as $P \wr H = G \rtimes_\varphi H$, where G is the n -fold direct product $\underbrace{P \times P \times \cdots \times P}_{n \text{ times}}$ and $\varphi : H \rightarrow \text{Aut}(G)$ is given by

$$\varphi_h(g_1, \dots, g_n) = (g_{h^{-1}(1)}, g_{h^{-1}(2)}, \dots, g_{h^{-1}(n)}). \quad (4)$$

The wreath product $P \wr H$ may be viewed as a subgroup of $\text{Sym}([m] \times [n])$ via the group homomorphism

$$\begin{aligned} \psi : G \rtimes_\varphi H &\rightarrow \text{Sym}([m] \times [n]) \\ (g, h) &\mapsto ((i, j) \mapsto (g_{h(j)}(i), h(j))) \end{aligned}$$

The action of $P \wr H$ on $[m] \times [n]$ given by ψ preserves the partition

$$[m] \times [n] = \bigcup_{k=1}^n [m] \times \{k\}.$$

This is the most important property of the wreath product.

At several points, we must distinguish between abstract groups and the way they act on sets. For instance, the standard action of S_4 on $[4]$ is not equivalent to the action of $S_4 \hookrightarrow S_6$ on the $6 = \binom{4}{2}$ unordered pairs in $[4]$. The latter group may also be described as $S_2 \wr S_3 \cap A_6$. It will reappear in Section 4.2 on homography estimation.

The permutation σ_γ in (3) is independent of the homotopy class of γ in U , from which one obtains a map from the fundamental group $\pi_1(U, z)$ into the symmetric group $\text{Sym}(X_z)$:

$$\begin{aligned} \rho : \pi_1(U, z) &\rightarrow \text{Sym}(X_z) \\ [\gamma] &\mapsto \sigma_\gamma. \end{aligned} \quad (5)$$

The map ρ in Equation 5 is known as the *monodromy representation*. The image of this map is a permutation group which acts transitively on the fiber X_z . It is called the *monodromy group*, and will be denoted by $\text{Mon}(X/Z; U, z)$, or $\text{Mon}(X/Z; z)$, or simply $\text{Mon}(X/Z)$. The latter notation, and our preferred terminology of *Galois/monodromy group*, is justified by Proposition 2.3. We write $\mathbb{C}(X)^{\text{gal}}/\mathbb{C}(Z)$ for the Galois closure of $\mathbb{C}(X)/\mathbb{C}(Z)$ and $\text{Gal}(X/Z)$ for its Galois group, which acts as a permutation group on the roots of some primitive element for the Galois extension $\mathbb{C}(X)^{\text{gal}}/\mathbb{C}(Z)$.

Proposition 2.3. Let $X \dashrightarrow Z$ be a branched cover with regular locus U , and fix a basepoint $z \in U$. Then the $\text{Gal}(X/Z)$ and $\text{Mon}(X/Z; U, z)$ are isomorphic as permutation groups. In particular, $\text{Mon}(X/Z; U, z)$ is independent of the choice of (U, z) .

A proof of Proposition 2.3 is given in [36]. In this proof, the Galois closure $\mathbb{C}(X)^{\text{gal}}/\mathbb{C}(Z)$ is identified as an extension of $\mathbb{C}(X)$ obtained by adjoining certain germs of functions around points in X_z . With this identification, it is then argued (using the Galois correspondence and analytic continuation) that the Galois and monodromy actions on X_z coincide.

Example 2.4. Consider

$$X = \{(x, z) \in \mathbb{C} \times \mathbb{C} \mid x^3 = z\}$$

as a degree-3 branched cover over $Z = \mathbb{C}$ given by $(x, z) \mapsto z$. A regular locus is the punctured complex line $U = \{z \mid z \neq 0\}$. The monodromy group $\text{Mon}(X/Z) \cong A_3 = C_3$ acts by cyclic

permutation of $X_z = \{(\sqrt[3]{z}, z), (\omega \sqrt[3]{z}, z), (\omega^2 \sqrt[3]{z}, z)\}$, where $\omega = \exp(2\pi i/3)$. Indeed, $\pi_1(U; z)$ is generated by the loop $\gamma(t) = e^{2\pi i t}$, which encircles the branch point $z = 0$ and induces the permutation σ_γ defined by

$$\sigma_\gamma = \begin{pmatrix} (\sqrt[3]{z}, z) & (\omega \sqrt[3]{z}, z) & (\omega^2 \sqrt[3]{z}, z) \\ (\omega \sqrt[3]{z}, z) & (\omega^2 \sqrt[3]{z}, z) & (\sqrt[3]{z}, z) \end{pmatrix}.$$

Remark 2.5. Although the branched covers we study are defined over the rational numbers, it is usually the geometric Galois group $\text{Gal}(\mathbb{C}(X)^{\text{gal}}/\mathbb{C}(Z))$ that interests us, rather than the arithmetic Galois group $\text{Gal}(\mathbb{Q}(X)^{\text{gal}}/\mathbb{Q}(Z))$. In general, the geometric Galois group is a subgroup of the arithmetic Galois group. The two can be different—the arithmetic Galois group for Example 2.4 is S_3 . However, we expect these two groups to be the same for most problems considered here.

Example 2.4 is atypical since X/Z is a *Galois cover*—in other words, the field extension $\mathbb{C}(X)/\mathbb{C}(Z)$ is Galois. Most branched covers encountered in computer vision are *not* Galois. Nevertheless, we can compute $\text{Mon}(X/Z)$ without computing the Galois closure via the numerical methods of Section 2.2. We consider a simple example of a non-Galois cover and its Galois closure below.

Example 2.6. Consider the parametric polynomial $p(x; a, b, c) = x^3 + ax^2 + bx + c$ and the corresponding degree-3 branched cover $X \rightarrow Z$ where $X = \{(x, a, b, c) \in \mathbb{C} \times \mathbb{C}^3 \mid p(x; a, b, c) = 0\}$ and $Z = \mathbb{C}^3$. A model of $\mathbb{C}(X)^{\text{gal}}$ in this case may be obtained by “cloning” the variable x three times and writing down Vieta’s equations for the roots of p :

$$X^{\text{gal}} = \{(x_1, x_2, x_3, a, b, c) \in \mathbb{C}^3 \times \mathbb{C}^3 \mid -(x_1 + x_2 + x_3) = a, (x_1x_2 + x_1x_3 + x_2x_3) = b, -x_1x_2x_3 = c\}.$$

The branched cover $X^{\text{gal}} \mapsto Z$ has degree 6. As abstract groups, both $X \rightarrow Z$ and $X^{\text{gal}} \rightarrow Z$ have the monodromy group S_3 . In the former case, it has the natural action on $[3]$, and in the latter case it is the left-regular representation of S_3 acting on itself. Proposition 2.16 implies that $\text{Mon}(X^{\text{gal}}/Z)$ is isomorphic as a permutation group to its centralizer in the ambient S_6 . This centralizer subgroup is the right-regular representation.

Definition 2.7. Two branched covers, $X_1 \dashrightarrow Z_1$ and $X_2 \dashrightarrow Z_2$, are *birationally equivalent* if there exist birational maps $X_1 \dashrightarrow X_2$ and $Z_1 \dashrightarrow Z_2$ such that the following diagram commutes:

$$\begin{array}{ccc} X_1 & \dashrightarrow & X_2 \\ \downarrow & & \downarrow \\ Z_1 & \dashrightarrow & Z_2. \end{array}$$

Proposition 2.8. The Galois/monodromy group of a branched cover is a birational invariant.

Proof. This follows easily from Proposition 2.3, since a birational equivalence of branched covers induces an isomorphism of field extensions. A more topological proof is also possible. Let us write $\Psi: X_1 \dashrightarrow X_2$, $\psi: Z_1 \dashrightarrow Z_2$ for the maps appearing in Definition 2.7. Then, for suitable regular loci $U_1 \subset Z_1, U_2 \subset Z_2$, there is an isomorphism $\text{Mon}(X_1/Z; U_1, z) \cong \text{Mon}(X_2/Z; U_2, \Psi(z))$ which may be defined by identifying, for each $\gamma: [0, 1] \rightarrow U_1$ based at z , the lifts $\tilde{\gamma}_1, \dots, \tilde{\gamma}_d$ in X_1 with the lifts $\Psi \circ \tilde{\gamma}_1, \dots, \Psi \circ \tilde{\gamma}_d$ in X_2 . $\square$

Example 2.9. Consider the regular branched cover $\mathrm{SO}_{\mathbb{C}}(3) \times \mathbb{P}(\mathbb{C}^{3 \times 1}) \rightarrow \mathcal{E}$ given by

$$(\mathbf{R}, \mathbf{t}) \mapsto [\mathbf{t}]_{\times} \mathbf{R}, \quad (6)$$

where for $\mathbf{t} = [t_1 : t_2 : t_3]$ we let

$$[\mathbf{t}]_{\times} = \begin{pmatrix} 0 & -t_3 & t_2 \\ t_3 & 0 & -t_1 \\ -t_2 & t_1 & 0 \end{pmatrix}$$

denote a 3×3 matrix that represents, up to scale, taking the cross product with $\mathbf{t}$, and $\mathcal{E}$ denotes the variety of *essential matrices*:

$$\mathcal{E} = \{\mathbf{E} \in \mathbb{P}(\mathbb{C}^{3 \times 3}) \mid \mathbf{E} \mathbf{E}^{\top} \mathbf{E} - \frac{1}{2} \mathrm{tr}(\mathbf{E} \mathbf{E}^{\top}) \mathbf{E} = 0\}. \quad (7)$$

A regular locus is given by $U \subset \mathcal{E}$ such that all $\mathbf{E} \in U$ have rank 2 and the kernel of $\mathbf{E}$ is not spanned by an isotropic vector where the twisted pair map is undefined. A birationally equivalent branched cover was constructed in [62, 92], where the authors construct moduli spaces obtained by letting the *absolute conic* [38] degenerate to a double line. Explicitly, the branched cover $X \rightarrow \mathcal{E}$ is given by

$$X = \{([a_0 : a_1 : a_2 : a_3], [b_0 : b_1 : b_2 : b_3]) \in \mathbb{P}^3 \times \mathbb{P}^3 \mid a_0 b_0 + a_1 b_1 + a_2 b_2 + a_3 b_3 = 0\}$$

$$([\mathbf{a}], [\mathbf{b}]) \mapsto \begin{pmatrix} a_0 b_0 - a_1 b_1 - a_2 b_2 + a_3 b_3 & a_1 b_0 + a_0 b_1 + a_3 b_2 + a_2 b_3 & a_2 b_0 - a_3 b_1 + a_0 b_2 - a_1 b_3 \\ a_1 b_0 + a_0 b_1 - a_3 b_2 - a_2 b_3 & -a_0 b_0 + a_1 b_1 - a_2 b_2 + a_3 b_3 & a_3 b_0 + a_2 b_1 + a_1 b_2 + a_0 b_3 \\ a_2 b_0 + a_3 b_1 + a_0 b_2 + a_1 b_3 & -a_3 b_0 + a_2 b_1 + a_1 b_2 - a_0 b_3 & -a_0 b_0 - a_1 b_1 + a_2 b_2 + a_3 b_3 \end{pmatrix},$$

and there exists a birational equivalence

$$\begin{array}{ccc} X & \dashrightarrow & \mathrm{SO}_{\mathbb{C}}(3) \times \mathbb{P}^2 \\ \downarrow & & \downarrow \\ \mathcal{E} & \longrightarrow & \mathcal{E}. \end{array}$$

where the bottom map is the identity. The top map may be given by

$$([\mathbf{a}], [\mathbf{b}]) \mapsto \left((a_3 I + [\mathbf{a}]_{\times})([\mathbf{a}]_{\times} - a_3 I)^{-1}, \begin{pmatrix} 2a_1 b_0 - 2a_0 b_1 + 2a_3 b_2 - 2a_2 b_3 \\ 2a_2 b_0 - 2a_3 b_1 - 2a_0 b_2 + 2a_1 b_3 \\ 2a_3 b_0 + 2a_2 b_1 - 2a_1 b_2 - 2a_0 b_3 \end{pmatrix} \right),$$

where now

$$[\mathbf{a}]_{\times} = \begin{pmatrix} 0 & a_0 & a_1 \\ -a_0 & 0 & a_2 \\ -a_1 & -a_2 & 0 \end{pmatrix}.$$

The map $X \dashrightarrow \mathrm{SO}_{\mathbb{C}}(3)$ is undefined for $([\mathbf{a}], [\mathbf{b}])$ such that $\mathbf{a}$ lies on the isotropic quadric $\|\mathbf{a}\|^2 = 0$ in $\mathbb{P}^3$. In [92], it was observed that a regular locus of $X \rightarrow \mathcal{E}$ is simply given by any $U \subset \mathcal{E}$ such that all $\mathbf{E} \in U$ have rank 2. Both branched covers have degree 2, and $\mathrm{Mon}(X/\mathcal{E})$ is the symmetric group S_2 acting on two letters. The map $X \dashrightarrow \mathrm{SO}_{\mathbb{C}}(3) \times \mathbb{P}^2$ is essentially Study's parametrization $\mathrm{SE}_{\mathbb{C}}(3)$ modulo scaling of translations, as may be seen by comparison with formulas [46, Eqs. 2.3, 2.4]. This parametrization is closely related to dual quaternions [46, Sec 2.4].

A *factorization* of a branched cover $X \dashrightarrow Z$ is a commutative diagram

$$\begin{array}{ccc} X & \dashrightarrow & Y \\ & \searrow & \downarrow \\ & & Z \end{array} \quad (8)$$

such that $X \dashrightarrow Y$ and $Y \dashrightarrow Z$ are branched covers. If $\deg(X/Y)$ and $\deg(Y/Z)$ are both strictly less than $\deg(X/Z)$, we say that the factorization is *proper* and that the branched cover $X \dashrightarrow Z$ is *decomposable*. Otherwise, $X \dashrightarrow Z$ is *indecomposable*. Proposition 2.10 implies that the decomposability of a branched cover can be determined from the Galois/monodromy group alone. We recall that a *block system* for the monodromy action $\text{Mon}(X/Z) \curvearrowright X_z = \{x_1 \dots x_d\}$, is a partition of $X_z = B_1 \cup \dots \cup B_k$, comprised of equally-sized blocks $B_1, \dots, B_k$, which is preserved in the sense that blocks are always mapped to blocks under the group action. The block systems associated to the action form a lattice under refinement, whose respective maximum and minimum elements are $\{X_z\}$ and $\{\{x_1\}, \dots, \{x_d\}\}$. If any other block systems exist, then $\text{Mon}(X/Z)$ is said to be *imprimitive*, and otherwise it is *primitive*.

Given a factorization (8), we have $\deg(X/Z) = \deg(X/Y) \deg(Y/Z)$, since a sufficiently generic point $z \in Z$ has $\deg(X/Z)$ preimages in X . It also has $\deg(Y/Z)$ preimages in Y , each of which have $\deg(X/Y)$ preimages in X . Moreover, there is a partition

$$X_z = X_{y_1} \cup \dots \cup X_{y_k} \quad (9)$$

with $k = \deg(Y/Z)$. The proof of Proposition 2.11 below shows that this is a block system for the monodromy action with blocks of size $\deg(X/Y)$. Conversely, imprimitivity implies decomposability.

Proposition 2.10. A branched cover is decomposable if and only if $\text{Mon}(X/Z)$ is imprimitive.

Proposition 2.10 dates back to the work of Ritt [80], who characterized the possible decompositions of branched covers $\mathbb{C} \ni x \mapsto p(x) \in \mathbb{C}$ given by a univariate polynomial p . A Galois-theoretic proof of Proposition 2.10 may be found, for instance, in [14]. If we know $\text{Mon}(X/Z)$, it is also possible to identify $\text{Mon}(X/Y)$ and $\text{Mon}(Y/Z)$ occurring in the factorization (8), as the next proposition shows.

Proposition 2.11. Consider a factorization of branched covers as in Equation (8). For fixed generic $z \in Z$, partition X_z as in Equation 9. The action $\text{Mon}(X/Z) \curvearrowright X_z$ induces two other group actions which are equivalent to the monodromy groups of the individual factors:

- 1) action on blocks: $\text{Mon}(X/Z) \curvearrowright \{X_{y_1}, \dots, X_{y_k}\}$, which is equivalent to $\text{Mon}(Y/Z)$.
- 2) action on a single block: $\text{Mon}(X/Z)_{X_y} \curvearrowright X_y$, where $\text{Mon}(X/Z)_{X_y}$ denotes the stabilizer of the set X_y under the action by $\text{Mon}(X/Z) \curvearrowright X_z$. This is equivalent to $\text{Mon}(X/Y)$, and thus independent of the choice $y \in Y_z$.

Proof. 1) For each $\sigma_\gamma \in \text{Mon}(X/Z)$, there is an induced permutation of the blocks:

$$\widetilde{\sigma}_\gamma = \begin{pmatrix} X_{y_1} & \cdots & X_{y_k} \\ \sigma_\gamma(X_{y_1}) & \cdots & \sigma_\gamma(X_{y_k}) \end{pmatrix}.$$

Indeed, suppose that $x, x' \in X_{y_i}$ are such that $\sigma_\gamma(x) \in X_{y_j}, \sigma_\gamma(x') \in X_{y_k}$, and consider the lift $\tilde{\gamma} : [0, 1] \rightarrow Y$ starting at y_i . We must have both $\tilde{\gamma}(1) = y_j$ and $\tilde{\gamma}(1) = y_k$. Hence $k = j$ by the unique path-lifting property applied to $Y \dashrightarrow Z$, showing that σ_γ preserves the partition into blocks. In this way we get a group homomorphism

$$\begin{aligned} \text{Mon}(X/Z) &\rightarrow \text{Sym}(\{X_{y_1}, \dots, X_{y_k}\}) \\ \sigma_\gamma &\mapsto \widetilde{\sigma}_\gamma, \end{aligned} \quad (10)$$

which represents the action of $\text{Mon}(X/Z)$ on the blocks. Now, there is also an injective group homomorphism

$$\begin{aligned} \text{Mon}(Y/Z) &\rightarrow \text{Sym}(\{X_{y_1}, \dots, X_{y_k}\}) \\ \tau_\gamma &\mapsto \begin{pmatrix} X_{y_1} & \cdots & X_{y_k} \\ X_{\tau_\gamma(y_1)} & \cdots & X_{\tau_\gamma(y_k)} \end{pmatrix} \end{aligned} \quad (11)$$

obtained by restricting the natural isomorphism $\text{Sym}(Y_z) \cong \text{Sym}(\{X_{y_1}, \dots, X_{y_k}\})$ that identifies a point $y_i \in Y_z$ with its corresponding block X_{y_i} . We wish to show that maps (10) and (11) have the same image. This follows easily if we restrict γ in both maps to be loops contained in a regular locus for $X \dashrightarrow Z$: the lifts of γ to Y (which, by our restriction, also lift to X) determine the corresponding permutation of blocks in X , and vice-versa. Indeed, we have $\sigma_\gamma(X_y) = X_{\tau_\gamma(y)}$ for any $y \in Y_z$. To see this, it is enough to show one set is contained in the other. A point $x \in \sigma_\gamma(X_y)$ is the endpoint of some lift of γ to X . The image of this lift in Y is itself a lift $\tilde{\gamma} : [0, 1] \rightarrow Y$ with $\tilde{\gamma}(0) = y$ —hence $\tilde{\gamma}(1) = \tau_\gamma(y)$, and the endpoint of our original lift x is in $X_{\tau_\gamma(y)}$.

2) The proof amounts to showing that a loop γ in Z lifts to a loop in Y if and only if $\sigma_\gamma \in \text{Mon}(X/Z)$ stabilizes each of the blocks. As in the previous part, this is only true if we consider loops in a suitably small regular locus $U \subset Z$. It suffices to take U contained in a regular locus for $X \dashrightarrow Z$ and whose preimage in Y is a regular locus for $X \dashrightarrow Y$. $\square$

Proposition 2.10 shows that an arbitrary branched cover $X \dashrightarrow Z$ factors as a composition of indecomposable branched covers:

$$X = Y_0 \dashrightarrow Y_1 \dashrightarrow \cdots \dashrightarrow Y_{k-1} \dashrightarrow Y_k = Z. \quad (12)$$

Such a factorization corresponds to a maximal chain in the lattice of block systems, and the associated degrees can be read off from the block sizes. Equivalently, for any $x \in X_z$, a maximal chain in the lattice of blocks of $\text{Mon}(X/Z)$ corresponds to a chain of subgroups that contain the stabilizer $\text{Mon}(X/Z)_x$ (cf. [81, proof of Theorem 9.15])

$$\text{Mon}(X/Z)_x = G_0 \subset G_1 \subset \cdots \subset G_{k-1} \subset G_k = \text{Mon}(X/Z),$$

and we have $\deg(Y_i/Y_{i+1}) = [G_{i+1} : G_i]$ for $i = 0, \dots, k-1$. The decomposition (12) is not unique. In fact, as Ritt already understood [80, p. 53], there are many examples where even the multi-set of degrees $\deg(Y_i/Y_{i+1})$ is not unique. See [34, Example 25] for one such explicit example.

Finally, we define and carefully study the deck transformations of a branched cover, of which the twisted pair symmetry from the introduction is a special case.

Definition 2.12. A birational equivalence from a branched cover to itself which fixes the base is called a *deck transformation*. Explicitly, for $f : X \dashrightarrow Z$ a deck transformation $\Psi : X \dashrightarrow X$ must satisfy $f \circ \Psi = f$ whenever both maps are defined. The deck transformations form a group under composition which acts on a generic fiber X_z . The deck transformation group can be naturally identified with the automorphisms of $\mathbb{C}(X)$ which fix $\mathbb{C}(Z)$, denoted $\text{Aut}(X/Z)$.

Analogously to decomposability, Proposition 2.16 shows that the existence of a nontrivial deck transformation can be decided from the Galois/monodromy group alone. This turns out to be stronger than decomposability in general. We learned of Proposition 2.16 from the sources [7, 18]. Since it seems less well-known outside of the literature on Galois/monodromy groups, we give a




Figure 3: Frontal view of the P3P problem: x_1, x_2, x_3 are unknown.

self-contained proof. In topology, a deck transformation of a covering map f can be any continuous function Ψ satisfying $f \circ \Psi = f$. Our proof of Proposition 2.16 reveals that, for a rational branched cover f with regular locus U , the deck transformations of $f|_{f^{-1}(U)}$ in the topological sense are always rational maps in the sense of Definition 2.12. Before giving the proof, we first consider three illustrative examples.

Example 2.13. Let $X = \mathbf{V}(x^2 + ax + b) \subset \mathbb{C}^3$, $Z = \mathbb{C}^2$ and $f : X \rightarrow Z$ be the degree-2 branched cover defined by coordinate projection $f(x, a, b) = (a, b)$. The deck transformation defined by $\Psi(x, a, b) = (-x - a, a, b)$ acts on a generic fiber $X_{(a,b)}$ by permuting the two roots of the quadratic equation $x^2 + ax + b = 0$.

Example 2.14. Ask et al. [6] define a polynomial system $F(\mathbf{x})$ with p -fold symmetry to be such that $F(\mathbf{x}) = 0$ implies $F(\omega \mathbf{x}) = 0$ whenever ω is a p -th root of unity. For example, the equations

$$\begin{aligned} f_{1,2} &= x_1^2 + x_2^2 - c_{1,2}x_1x_2 - d_{1,2}^2 \\ f_{1,3} &= x_1^2 + x_3^2 - c_{1,3}x_1x_3 - d_{1,3}^2 \\ f_{2,3} &= x_2^2 + x_3^2 - c_{2,3}x_2x_3 - d_{2,3}^2 \end{aligned} \tag{13}$$

have a 2-fold sign symmetry: $(x_1, x_2, x_3) \mapsto (-x_1, -x_2, -x_3)$. These equations define the famous Perspective-3-Point problem or *P3P problem*. Here, each $c_{i,j}$ is equal to $2 \cos \theta_{i,j}$ as in Figure 3. Letting X denote the vanishing locus of (13) in $\mathbb{C}^9$, the coordinate projection onto the space of knowns $\mathbb{C}^6$ is a branched cover with a deck transformation given by the sign-symmetry. We will return to the P3P problem in Section 3.

Ask et al. [6] develop algorithms for detecting and exploiting *partial* p -fold symmetries (occurring in only some subset of the variables) in the automatic generation of polynomial solvers. These methods were generalized by Larsson and Åström [55] to the case of *weighted* partial- p fold symmetries. In general, a branched cover with a weighted partial- p fold symmetry will have a deck transformation of order p , degree ep for some integer e , and its Galois/monodromy group will be a subgroup of $C_p \wr S_e$.

Example 2.15. Example 2.9 contains two interesting *Galois covers*—the Galois/monodromy group and the deck transformation group are isomorphic. For $\mathrm{SO}_{\mathbb{C}}(3) \times \mathbb{P}^2 \dashrightarrow \mathcal{E}$, the action on the fiber applies the twisted pair map as in (2). For $X \rightarrow \mathcal{E}$, the action swaps coordinates $([\mathbf{a}], [\mathbf{b}]) \mapsto ([\mathbf{b}], [\mathbf{a}])$.

Proposition 2.16. Let $X \dashrightarrow Z$ be a branched cover and fix generic $z \in Z$. We may identify the deck transformation group with a subgroup of $Sym(X_z)$ by restricting functions to X_z . This permutation group is *equal* to the centralizer of $\text{Mon}(X/Z)$ in $Sym(X_z)$.

Proof. We abbreviate the deck transformation group and centralizer subgroup by D and C , respectively. We define a map between these groups as follows:

$$\begin{aligned} \varphi : D &\rightarrow Sym(X_z) \\ \Psi &\mapsto \begin{pmatrix} x_1 & \cdots & x_d \\ \Psi(x_1) & \cdots & \Psi(x_d) \end{pmatrix} \end{aligned}$$

To prove Proposition 2.16, we verify the following properties of φ :

- 1) φ is a group homomorphism.
- 2) φ is injective.
- 3) The image of φ is contained in C .
- 4) C is contained in the image of φ —more explicitly, for all $\sigma \in C$ there exists a deck transformation $\Psi_\sigma \in D$ whose restriction to the fiber X_z equals the permutation σ .

Property 1) is straightforward. Properties 2) and 3) both follow from the unique path-lifting property. For instance, if $\Psi(x_i) = x_i$ for $i = 1, \dots, d$, then a generic $x \in X$ will be the endpoint of the lift $\tilde{\gamma}$ of some path γ in Z based at z —if $x_i = \Psi(x_i)$ is the initial point of this lift, then we must have $\Psi \circ \tilde{\gamma} = \tilde{\gamma}$, so that in particular $\Psi(x) = \Psi \circ \tilde{\gamma}(1) = \tilde{\gamma}(1) = x$. This gives Property 2). The proof of Property 3) is very similar, and may also be found, for instance, in [18, Proposition 1.3].

It remains to show Property 4). We do so by first constructing a map $\Psi_\sigma : X \dashrightarrow X$ pointwise via lifting paths. The argument is analogous to the proof in [40, Proposition 1.39]. Fix $x_0 \in X_z$. For generic $x \in X$, there exists a path $\alpha_x : [0, 1] \rightarrow X$ from x_0 to x whose image in Z is contained in a regular locus for $X \dashrightarrow Z$. Define $\overline{\alpha_x}$ to be the lift based at $\sigma(x_0)$ whose image in Z coincides with the image of α_x . We define

$$\begin{aligned} \Psi_\sigma : X &\dashrightarrow X \\ x &\mapsto \overline{\alpha_x}(1). \end{aligned}$$

First of all, we must show that Ψ_σ is well-defined. This means that for any other path β_x from x_0 to x we must have $\overline{\beta_x}(1) = \overline{\alpha_x}(1)$. We refer the reader to Figure 4 to more easily follow the argument. Consider the loop γ based at $f(x_0)$ in Z obtained by concatenating $\overleftarrow{f \circ \beta_x}$ (the reverse of the path $f \circ \beta_x$) with $f \circ \alpha_x$. Since σ and σ_γ commute, we have that

$$\begin{aligned} \sigma_\gamma(\sigma(x_0)) &= \sigma(\sigma_\gamma(x_0)) \\ &= \sigma(x_0). \end{aligned}$$

Thus σ_γ fixes $\sigma(x_0)$, and it follows that the lift $\tilde{\gamma}$ based at $\sigma(x_0)$ is a loop. This implies that $\overline{\alpha_x}$ and $\overline{\beta_x}$ have the same terminal point $\Psi_\sigma(x)$, proving well-definedness.

Consider now an arbitrary $x \in X_z$. Note that in this case $f \circ \alpha_x$ is a loop. We calculate

$$\begin{aligned} \sigma(x) &= \sigma(\sigma_{f \circ \alpha_x}(x_0)) \\ &= \sigma_{f \circ \alpha_x}(\sigma(x_0)) \\ &= \Psi_\sigma(x). \end{aligned}$$




Figure 4: Construction and well-definedness of Ψ_σ .

Thus, restricting Ψ_σ to X_z yields the permutation σ . Moreover, by definition of Ψ_σ we have that $f \circ \Psi_\sigma = f$ on the locus of points where both maps are defined. It remains to show that Ψ_σ is a rational map, since then it will also follow that $\Psi_{\sigma^{-1}}$ is a rational inverse.

First we note that, in a suitably small neighborhood of any generic point $x \in X$, we can write $\Psi_\sigma = g_x \circ f$, where g_x is a holomorphic local inverse of f . In more detail, if f and Ψ_σ are both defined at x , then there exists a neighborhood $U_x \ni x$ (in the complex topology) where both f and Ψ are defined and such that $\Psi_\sigma(U_x) = U_{\Psi_\sigma(x)}$ is a neighborhood containing $\Psi_\sigma(x)$. Shrinking these neighborhoods further as necessary, we may apply the inverse function theorem (in its holomorphic form or, equivalently, the holomorphic implicit function theorem as in [45, Theorem 2.1.2]) to the map $f|_{U_{\Psi_\sigma(x)}} : U_{\Psi_\sigma} \rightarrow U_z$, where U is some neighborhood of z . This theorem gives the needed local inverse $g_x : U_z \rightarrow U_{\Psi_\sigma(x)}$.

An expression of the form $\Psi_\sigma = g_x \circ f$ exists for any x in some Zariski open subset of X . It follows that Ψ_σ is a meromorphic map from X to itself—in other words, it is holomorphic after restricting to a Zariski-open $U \subset X$. To finish the proof, we may use the well-known fact that all meromorphic maps between projective varieties are rational.¹ Indeed, if X and Z are quasiprojective varieties, we may replace them with their projective closures $\overline{X}, \overline{Z}$, to get a birationally equivalent $\overline{X} \dashrightarrow \overline{Z}$. See also [85, Theorem 6.1.4], which states a general GAGA principle for branched covers of not-necessarily projective varieties. $\square$

Proposition 2.17. A branched cover $X \dashrightarrow Z$ of degree d with a nontrivial deck transformation Ψ is either decomposable or its Galois/monodromy group is cyclic of order d . In the latter case, $\text{Mon}(X/Z)$ is imprimitive precisely when d is composite.

Proof. Partition X_z into the orbits under repeated application of Ψ . This partition is preserved under

¹We recall the words of Mumford [70, Ch. 4]: “This should be viewed as a generalization of the old result that the only everywhere meromorphic functions on $\mathbb{C} \cup \{\infty\}$ are rational functions.”

the monodromy action. Thus, the Galois group is imprimitive if this partition is nontrivial. Otherwise, the action of Ψ on X_z generates a cyclic group $C_d \subset S_d$. Letting $\text{Cent}(\cdot)$ denote the centralizer in S_d , we have $C_d \subset \text{Cent}(\text{Mon}(X/Z))$, which holds if and only if $\text{Mon}(X/Z) \subset \text{Cent}(C_d) = C_d$. Since $\text{Mon}(X/Z)$ is transitive, we must have $\text{Mon}(X/Z) = C_d$. $\square$

In general, a decomposable branched cover need not have any deck transformations. However, a converse to Proposition 2.17 does hold in a special case frequently encountered in practice.

Proposition 2.18. $X \dashrightarrow Z$ has a deck transformation of order 2 if and only if $\text{Mon}(X/Z)$ has a block of size 2.

Proof. $\Rightarrow$ As in Proposition 2.17. $\Leftarrow$ Proposition 2.10 gives a factorization such that $\mathbb{C}(X)/\mathbb{C}(Y)$ is a degree 2 extension, and thus always Galois. $\square$

2.2 Numerically computing Galois/monodromy groups

In this paper, our main interest is in minimal problems. These typically give rise to branched covers of the form $X \dashrightarrow \mathbb{C}^m$, at least up to birational equivalence. We again emphasize that being a branched cover implies $\dim X = m$ and, for *generic measurements* $z \in \mathbb{C}^m$, that the fiber X_z is finite. This is essentially the definition of a minimal problem used in [20, 23], where a problem is minimal if and only if its joint camera map is a branched cover. For $X \subset \mathbb{C}^n$, coordinate projection from the graph $\text{Gr}(f) \rightarrow \mathbb{C}^m$ is a branched cover which is birationally equivalent to $X \dashrightarrow \mathbb{C}^m$. Proposition 2.8 implies that these branched covers' Galois/monodromy groups are isomorphic as permutation groups. Thus, in practice it is usually enough to consider the case where X is a subvariety of $\mathbb{C}^n \times \mathbb{C}^m$, and the map $X \rightarrow \mathbb{C}^m$ is coordinate projection. In this case, the ideal $\mathcal{I}_X$ of all polynomials in $\mathbb{C}[x_1, \dots, x_n, z_1, \dots, z_m]$ that vanish on X is prime.

However, one advantage of computing Galois/monodromy groups numerically is that we do not need to know all generators of $\mathcal{I}_X$. All that is really needed are

- 1) the ability to sample a generic point $(x^*, z^*) \in X$, and
- 2) a set of n equations vanishing on X ,

$$F(x; z) = F(x_1, \dots, x_n; z_1, \dots, z_m) = \begin{bmatrix} f_1(x_1, \dots, x_n; z_1, \dots, z_m) \\ \vdots \\ f_n(x_1, \dots, x_n; z_1, \dots, z_m) \end{bmatrix}, \quad (14)$$

such that the Jacobian $d_x F(x^*; z^*)$ is an invertible $n \times n$ matrix for generic $(x^*, z^*) \in X$. We say that equations (14) form a *well-constrained system* for the branched cover $X \rightarrow \mathbb{C}^m$.

In equation (14), $x_1, \dots, x_n$ are usually called the *variables* and $z_1, \dots, z_m$ are usually called the *parameters*. For generic $(x^*, z^*) \in X$ and generic $z \in \mathbb{C}^m$, the path $\alpha: [0, 1] \rightarrow \mathbb{C}^m$ defined by the straight-line segment $\alpha(t) = (1-t) \cdot z^* + t \cdot z$ will have a lift $\tilde{\alpha}(t)$ to X based at $\tilde{\alpha}(0) = (x^*, z^*)$. By genericity, the lifted path $\tilde{\alpha}: [0, 1] \rightarrow X$ will not intersect the subvariety of X where $d_x F$ is singular for any $t \in [0, 1]$ —see [89, Lemma 7.12]. This implies that $\tilde{\alpha}$ can be numerically approximated by applying numerical continuation methods to the *parameter homotopy*

$$H(x, t) = F(x; \alpha(t)) = 0, \quad (15)$$

which connects known solutions of the *start system* $H(x, 0) = F(x; z^*) = 0$ to solutions of the *target system*² $H(x, 1) = F(x; z) = 0$. The solution curves $x(t)$ satisfying $H(x(t), t) = 0$ and the initial condition $x(0) = x^*$ are such that $(x(t), \alpha(t))$ stays on the irreducible variety X with probability-one, and hence $\tilde{\alpha}(t) = (x(t), \alpha(t))$. Thus, although the variety X need not be a complete intersection, appropriate use of a well-constrained system enables us to compute solutions on X using the same number of equations as unknowns. This fits into an established paradigm of numerical algebraic geometry where overdetermined parameterized polynomial systems may be solved by reduction to a well-constrained system (see also [12, Sec 6.4], [41].) In practice, our numerical approximations to $\tilde{\alpha}(t)$ remain “close” to X with some probability that depends on the conditioning and the implementation of the numerical methods.

By numerically continuing solutions along some path $\beta(t)$ from z to z^* , and then along some other path $\alpha(t)$ from z^* to z , the concatenated path $\gamma = \beta * \alpha$ is a loop based at z which induces a monodromy permutation $\sigma_\gamma \in \text{Mon}(X/\mathbb{C}^m; z)$. This simple observation motivates numerous applications of monodromy in numerical algebraic geometry: computing the fibers X_z (addressed in [21, 64]) computing the Galois/monodromy group (addressed in [43, 61]), and in additional applications ranging from numerical irreducible decomposition [88] to kinematics [42]. To compute $\text{Mon}(X/\mathbb{C}^m)$ in this paper, we simply generate some number of loops $\gamma_1, \dots, \gamma_k$ in $\mathbb{C}^m$, with k ranging from 4 (usually sufficient when $\text{Mon}(X/\mathbb{C}^m)$ is full-symmetric) to as large as 50. This adds additional uncertainty to our numerical computations, since *a priori* we only know that $\langle \sigma_{\gamma_1}, \dots, \sigma_{\gamma_k} \rangle$ is a subgroup of $\text{Mon}(X/\mathbb{C}^m; z)$. In principle, this additional uncertainty could be avoided by a more computation-heavy approach like the branch-point method in [43]. Nevertheless, we feel reasonably confident in the Galois/monodromy group computations which we have reported, which have been validated through repeated runs and analyzing decompositions in several of the imprimitive cases.

Algorithm 1 Computing the Galois/monodromy group

Input: A parameterized system $F(x; z)$, well-constrained for some branched cover, that is polynomial in $x \in \mathbb{C}^n$ and $z \in \mathbb{C}^m$; a generic basepoint parameter value $z^* \in \mathbb{C}^m$.

Output: The Galois/monodromy group $\text{Mon}(X/Z)$.

- 1: Collect $d = \deg(X/Z)$ nonsingular isolated solutions at the basepoint z^* such that $F(x; z^*) = 0$, and fix an ordering of these solutions: $x^{(1)}, \dots, x^{(d)}$.
 - 2: Collect permutations $\sigma_{\gamma_1}, \dots, \sigma_{\gamma_k}$ generated from tracking parameter homotopies representing loops starting and ending at z^* .
 - 3: The permutation group $\langle \sigma_{\gamma_1}, \dots, \sigma_{\gamma_k} \rangle$ is a subgroup of $\text{Mon}(X/Z)$. For $k \gg 0$, we expect $\langle \sigma_{\gamma_1}, \dots, \sigma_{\gamma_k} \rangle = \text{Mon}(X/Z)$.
 - 4: **Return** $\langle \sigma_{\gamma_1}, \dots, \sigma_{\gamma_k} \rangle$.
-

²This convention is chosen to agree with the notation of the previous section. We note that this is the opposite of the common convention of placing start systems at $t = 1$ and target systems at $t = 0$ which, although mathematically equivalent, is more natural from the numerical point of view.

3 Absolute pose problems

In this section, we apply the mathematical framework of the previous section to absolute pose problems involving combinations of point/line features appearing in work of Ramalingam et al. [78] Absolute camera pose estimation is one of the main problems of computer vision [4, 28, 35, 59, 75, 79, 91, 94]. Although the problems considered here are of low degree, computing the Galois/monodromy groups yields new insights which might be applied to building better solvers for these problems.

We begin formulating these problems in the language of branched covers. Our general task is to determine a calibrated camera matrix $[\mathbf{R} \mid \mathbf{t}]$ from correspondence data between the scene and images. We let p and l be the numbers of point-point and line-line correspondences, respectively, between 3D and 2D. The total space of our branched cover is

$$X_{p,l} = (\mathbb{P}^3)^p \times (\mathbb{G}_{1,3})^l \times \mathrm{SE}_{\mathbb{C}}(3)$$

where $\mathbb{G}_{1,3}$ denotes the Grassmannian of lines in $\mathbb{P}^3$. The base space equals

$$Z_{p,l} = (\mathbb{P}^3)^p \times (\mathbb{G}_{1,3})^l \times (\mathbb{P}^2)^p \times (\mathbb{G}_{1,2})^l,$$

where $\mathbb{G}_{1,2}$ denotes the Grassmannian of lines in $\mathbb{P}^2$ and $f_{p,l}: X_{p,l} \dashrightarrow Z_{p,l}$ is the map that “takes pictures”:

$$\begin{aligned} & \left(X_1, \dots, X_p, \overline{L_1 L'_1}, \dots, \overline{L_l L'_l}, [\mathbf{R} \mid \mathbf{t}] \right) \mapsto \\ & \left(X_1, \dots, X_p, \overline{L_1 L'_1}, \dots, \overline{L_l L'_l}, [\mathbf{R} \mid \mathbf{t}] X_1, \dots, [\mathbf{R} \mid \mathbf{t}] X_p, \right. \\ & \quad \left. \overline{[\mathbf{R} \mid \mathbf{t}] L_1 [\mathbf{R} \mid \mathbf{t}] L'_1}, \dots, \overline{[\mathbf{R} \mid \mathbf{t}] L_l [\mathbf{R} \mid \mathbf{t}] L'_l} \right) \end{aligned}$$

(here $\overline{LL'}$ is the line spanned by two world points L and L' .) Counting dimensions gives $\dim X_{p,l} = 3(p+l) + 6$ and $\dim Z_{p,l} = 5(p+l)$. Equating the two, we see that the only possibilities are $(p,l) = (3,0), (2,1), (1,2), (0,3)$. The first case corresponds to the P3P problem.

Result 3.1. The full list of Galois/monodromy groups of branched covers $f_{p,l}$ is as follows:

$$\begin{aligned} \mathrm{Mon}(X_{3,0}/Z_{3,0}) &\cong S_2 \wr S_4 \cap A_8 \hookrightarrow S_8 \\ \mathrm{Mon}(X_{2,1}/Z_{2,1}) &\cong S_2 \wr S_2 \hookrightarrow S_4 \\ \mathrm{Mon}(X_{1,2}/Z_{1,2}) &\cong S_2 \wr S_4 \cap A_8 \hookrightarrow S_8 \\ \mathrm{Mon}(X_{0,3}/Z_{0,3}) &\cong S_8. \end{aligned}$$

We interpret Result 3.1 in two separate subsections, corresponding to the “unmixed cases” $(p,l) \in \{(3,0), (0,3)\}$ and the more interesting “mixed cases” $(p,l) \in \{(2,1), (1,2)\}$. We note the respective degrees 8, 4, 8, 8 agree with those reported in [78], in which these problems were formulated using different systems of equations. The systems of equations defining the parameter homotopies used for Result 3.1 were constructed as follows:

- Points in the world are represented by 4×1 matrices $X_1, \dots, X_p$.
- Points in the image are represented by 3×1 matrices $x_1, \dots, x_p$.
- Lines in the world are represented as kernels of 2×4 matrices $[\mathbf{N}_1 \mid \mathbf{N}'_1]^\top, \dots, [\mathbf{N}_l \mid \mathbf{N}'_l]^\top$.

- Lines in the image are represented as kernels of 1×3 matrices $\mathbf{n}_1^\top, \dots, \mathbf{n}_l^\top$.
- We enforce rank constraints by the vanishing of maximal minors of certain matrices:
 - point-to-point: $\text{rk} \left([\mathbf{R} \mid \mathbf{t}] X_i \mid x_i \right) \leq 1$ for $i = 1, \dots, p$
 - line-to-line: $\text{rk} \left(\mathbf{N}_i \mid \mathbf{N}'_i \mid [\mathbf{R} \mid \mathbf{t}]^\top \mathbf{n}_i \right) \leq 2$ for $i = 1, \dots, l$
- Any square subsystem of these maximal minors whose Jacobian has full rank gives a well-constrained system in the sense of Section 2.2 (the variables being $[\mathbf{R} \mid \mathbf{t}]$, and the parameters being all points and lines.) Although the square subsystem may have excess solutions, they are not in the same orbit as the geometrically relevant solutions under the monodromy action. Alternatively, we can get a well-constrained system by randomization as in [12, Sec 6.4]. The monodromy group does not depend on the choice of well-constrained system.

3.1 The unmixed cases: $(p, l) = (3, 0), (0, 3)$

The case $(p, l) = (3, 0)$ reduces to solving the P3P problem as formulated in Equations (13). The literature on this problem is vast, and the earliest work [33] pre-dates the field of computer vision by more than a century. The degree of this problem is 8 and the Galois/monodromy group is a subgroup of $S_2 \wr S_4$ due to the sign symmetry. The imprimitivity of $\text{Mon}(X_{3,0}/Z_{3,0})$ is due to the sign symmetry noted in Example 2.14. This sign symmetry can also be understood in terms of the monomial supports of Equations (13), which span an index-2 sublattice of $\mathbb{Z}^3$. In the terminology of Brysiewicz et al. [14], Equations (13) define a *lacunary* polynomial system. However, the Galois/monodromy group for this system cannot be read off from the monomial supports alone. In the setting of [14], we would consider the family of all systems with the same monomial supports as in (13)

$$\begin{aligned} h_{1,2} &= Ax_1^2 + Bx_2^2 + Cx_1x_2 + D \\ h_{1,3} &= Ex_1^2 + Fx_3^2 + Gx_1x_3 + H \\ h_{2,3} &= Ix_2^2 + Jx_3^2 + Kx_2x_3 + L \end{aligned} \tag{16}$$

This gives a branched cover $X_h \rightarrow \mathbb{C}^{12}$ where $X_h = V(h_{1,2}, h_{1,3}, h_{2,3}) \subset \mathbb{C}^3 \times \mathbb{C}^{12}$. On the other hand, for P3P the natural branched cover is $X_f \rightarrow \mathbb{C}^6$, where $X_f \subset \mathbb{C}^3 \times \mathbb{C}^6$. We find numerically that $\text{Mon}(X_h/\mathbb{C}^{12})$ is the full wreath product $S_2 \wr S_4$, whereas our numerical experiments suggest that the Galois/monodromy group for P3P is the *proper subgroup* $S_2 \wr S_4 \cap A_8$.

The four absolute pose problems appearing in Result 3.1 have small enough degree to allow the **independent** verification of their Galois/monodromy groups with symbolic computation. We illustrate one possible route for this verification for P3P problem $X_{3,0}/Z_{3,0}$. Consider $I = \langle f_{1,2}, f_{1,3}, f_{2,3} \rangle$ as an ideal in a polynomial ring $\mathbb{F}[x_1, x_2, x_3]$ whose coefficient field is $\mathbb{F} = \mathbb{C}(Z) = \mathbb{C}(\vec{c}, \vec{d})$. The dimension and degree of I are 0 and 8. We can compute a lexicographic Gröbner basis for I with $x_1 > x_2 > x_3$ in a matter of seconds using the FGLM Gröbner basis conversion algorithm [27], implemented for Macaulay2 [32] in the package FGLM [74]. The Gröbner basis $G = \{g_1, g_2, g_3\}$ has the form predicted by the Shape lemma:

$$\begin{aligned} g_1(x_1, x_2, x_3) &= x_1 + r_1(\vec{c}, \vec{d}) x_3 \\ g_2(x_2, x_3) &= x_2 + r_2(\vec{c}, \vec{d}) x_3 \\ g_3(x_3) &= x_3^8 + A(\vec{c}, \vec{d}) x_3^6 + B(\vec{c}, \vec{d}) x_3^4 + C(\vec{c}, \vec{d}) x_3^2 + D(\vec{c}, \vec{d}), \end{aligned}$$

for particular rational functions $r_1, r_2, A, B, C, D \in \mathbb{F}$. We see that x_3 is a primitive element for the extension $\mathbb{C}(X)/\mathbb{C}(Z)$. To verify that $\text{Mon}(X/Z) \cong \text{Gal}(X/Z)$ is contained in A_8 , it suffices to show that the discriminant of g_3 is square. For arbitrary coefficients (A, B, C, D) , the discriminant of $x_3^8 + Ax_3^6 + Bx_3^4 + Cx_3^2 + D$ is the product of D and a square. For P3P, $D(\vec{c}, \vec{d})$ is also a square.

Factorizations of P3P are also classical: from Equations (13), we have

$$\begin{aligned} y_1(1 + y_2^2 - c_{1,2}y_2) - d_{1,2}^2 &= 0 \\ y_1(1 + y_3^2 - c_{1,3}y_3) - d_{1,3}^2 &= 0 \\ y_1(y_2^2 + y_3^2 - c_{2,3}y_2y_3) - d_{2,3}^2 &= 0 \end{aligned} \tag{17}$$

where $y_1 = x_1^2$, $y_2 = x_2/x_1$, $y_3 = x_3/x_1$ are separating invariants [47] for the action of the deck transformation group. We see that even when $X \rightarrow Z$ is regular in the definition of a factorization (8), the maps $X \dashrightarrow Y$ and $Y \dashrightarrow Z$ need not be.

On the other hand, the branched cover $X_{0,3} \dashrightarrow Z_{0,3}$ for absolute pose from 3 lines is indecomposable, since its Galois/monodromy group S_8 acts primitively on the set of 8 solutions. Thus, any algebraic algorithm for solving this problem must be capable of computing the roots of a polynomial of degree 8 or higher.

3.2 The mixed cases: $(p, l) = (2, 1), (1, 2)$

Proposition 2.16 shows that each of the mixed cases has a nontrivial deck transformation group: we have that $\text{Aut}(X_{2,1}/Z_{2,1}) \cong \text{Aut}(X_{1,2}/Z_{1,2}) \cong C_2$. Using the rank constraints described above, we were able to observe numerically that solutions in the same block for both of these mixed cases differed by a reflection. These deck transformations take on a particularly simple form after changing coordinates as in [78].

For the case $(p, l) = (2, 1)$, the formulation [78, Equations 4,5] makes use of a clever choice of reference frames to get equations

$$\begin{aligned} \mathbf{A} \mathbf{X} - \mathbf{b} &= 0 \\ R_{1,1}^2 + R_{2,1}^2 + R_{3,1}^2 - 1 &= 0 \\ R_{2,1}^2 + R_{2,2}^2 + R_{2,3}^2 - 1 &= 0 \end{aligned} \tag{18}$$

where $\mathbf{A}$ and $\mathbf{b}$ are 6×8 and 8×1 matrices which are rational functions of the data $z \in Z_{2,1}$, and

$$\mathbf{X} = [\vec{R} \mid \mathbf{t}] = [R_{1,1}, R_{2,1}, R_{3,1}, R_{2,2}, R_{2,3}, t_1, t_2, t_3]^\top$$

is a vector of indeterminates using only the first column and last row from the unknown rotation $\mathbf{R}$. With this formulation, we may easily verify that there is a nontrivial deck transformation given by

$$(\vec{R}, \mathbf{t}, z) \mapsto (-\vec{R}, -\mathbf{t} - 2\mathbf{e}_3, z). \tag{19}$$

(Here $\mathbf{e}_3 \in \mathbb{R}^3$ is the third standard basis vector.) Likewise, for the $(p, l) = (1, 2)$ case, there is a similar formulation given by [78, Equations 7,8]. In this formulation, $\mathbf{A}$ and $\mathbf{b}$ are now 6×9 and 9×1 matrices depending rationally on $z \in Z_{1,2}$, and

$$\mathbf{X} = [R_{1,1}, R_{1,2}, R_{3,1}, R_{2,2}, R_{2,3}, t_1, t_2, t_3]^\top$$

is a vector of indeterminates using only the first two rows of $\mathbf{R}$. Here, once again, there is a deck transformation of the form (19).

For both of the mixed cases, the form of the deck transformation (19) suggests a strategy for solving system (18) by sequentially solving systems of smaller degree; eliminate the variable t_3 , and solve for the remaining variables up to the sign-symmetry. With hindsight, the symmetries generating $\text{Aut}(X_{2,1}/Z_{2,1})$ and $\text{Aut}(X_{1,2}/Z_{1,2})$ are quite easy to verify. However, we stress that computing the Galois/monodromy group was what led us to discover them.

4 Relative pose problems

Our interest in Galois/monodromy groups of minimal problems began when we computed, using Bertini [11], that the Galois/monodromy group of the five-point problem was $S_2 \wr S_{10} \cap A_{20}$. Much like P3P, we had initially expected the full wreath product. Since then, we have computed Galois/monodromy groups of many minimal problems using Bertini and the Macaulay2 package `MonodromySolver` [22]. Overall, we agree with the assessment of Esterov and Lang that Galois/monodromy groups of structured polynomial systems are “*unexpectedly rich*” [25].

Many of the problems we considered appeared previously in [20, Table 1]. In that work, branched covers were represented by pictograms. For instance, the five-point problem was denoted by . The problems  and  were introduced in [49, 26], respectively. Our computations that $\text{Mon}(\text{pictogram}) \cong S_{216}$ and $\text{Mon}(\text{pictogram}) \cong S_{312}$ show that the homotopy solvers developed in [26] are *optimal* in the sense of tracking the fewest paths possible.

In future work, we hope to investigate the large number of minimal problems involving partial correspondences [23], which may arise due to missing data or relaxation of some overconstrained, non-minimal problem. For example, the problem  appearing in [23] assumes we are given the data three complete point correspondences across three views, plus two partial “point–point–nothing” correspondences. The task is to recover the relative pose of three calibrated cameras that see these data. The Galois/monodromy group of this problem is imprimitive, as its branched cover is essentially a composition of the P3P and five-point problem branched covers. Unlike the majority of problems studied here, this composite minimal problem has an intermediate field $\mathbb{C}(Z) \subsetneq \mathbb{C}(Y) \subsetneq \mathbb{C}(X)$ which is not the fixed field of some subgroup of $\text{Aut}(X/Z)$.

Result 4.1. Among all minimal problems of degree $< 1,000$ appearing in [20, Table 1], all have either an imprimitive or full symmetric Galois/monodromy group. The imprimitive cases are:

$$\begin{aligned} \text{Mon}(\text{pictogram}) &\cong (C_2)^2 \rtimes (S_2 \wr S_3 \cap A_6) \hookrightarrow S_{12} \\ \text{Mon}(\text{pictogram}) &\cong S_2 \wr S_8 \cap A_{16} \hookrightarrow S_{16} \\ \text{Mon}(\text{pictogram}) &\cong S_2 \wr S_{10} \cap A_{20} \hookrightarrow S_{20} \\ \text{Mon}(\text{pictogram}) &\cong S_2 \wr (S_2 \wr S_{16} \cap A_{32}) \cap A_{64} \hookrightarrow S_{64} \\ \text{Mon}(\text{pictogram}) &\cong (C_2)^4 \rtimes \left((C_2)^4 \rtimes (S_2 \wr (S_2 \wr S_4)) \right) \hookrightarrow S_{64} \\ \text{Mon}(\text{pictogram}) &\cong (C_2)^2 \rtimes (C_2^2 \rtimes (S_2 \wr (S_2 \wr S_2 \cap A_4) \cap A_8)) \hookrightarrow S_{32}. \end{aligned}$$

For the sake of uniformity, we have used the semidirect product $\rtimes$ to indicate subgroups of an appropriate wreath product. Thus, for instance, for $\text{Mon}(\text{pictogram})$, the outermost $(C_2)^2$ should be regarded as a subgroup of $(S_2)^{16}$, and the innermost as a subgroup of $(S_2)^8$. Much to our surprise, the group $\text{Mon}(\text{pictogram})$ turns out to be solvable.

Clearly there are many minimal problems waiting to be decomposed. With increasingly efficient and user-friendly homotopy continuation software like Bertini, `HomotopyContinuation.jl` [13], and

various numerical packages for Macaulay2 (summarized in [60]), we see no obstacles to computing even more Galois/monodromy groups of interest to computer vision in the future.

In the remainder of this section, we give particular attention to three problems appearing in Result 4.1. The first is the five-point problem $\begin{smallmatrix} \bullet & \bullet \\ \bullet & \bullet \end{smallmatrix}$. The second is the five-point problem for data which lie in a “V”-shape $\begin{smallmatrix} \bullet & \bullet \\ \bullet & \bullet \end{smallmatrix}$. Decomposing the associated branched cover reveals the classical planar calibrated homography problem. The third, denoted $\begin{smallmatrix} \bullet & \bullet \\ \bullet & \bullet \end{smallmatrix}$, is a special case of the notorious four-points-in-three-views problem [73] in which three of the corresponding points are collinear.

4.1 Five-point relative pose

We now return to the five-point problem, retaining the notation from Examples 1.1 and 2.9. It is well-known fact that the branched cover $X \rightarrow Z$ is decomposable. The intermediate variety is

$$Y = \{(\mathbf{E}, (\mathbf{x}_1, \dots, \mathbf{y}_5)) \in \mathcal{E} \times Z \mid \mathbf{y}_i^\top \mathbf{E} \mathbf{x}_i = 0, i = 1, \dots, 5\}.$$

The equations defining Y give the formulation of the five-point problem studied in several seminal works [19, 51, 66], and used in state of the art five-point solvers such as Nistér’s [72]. The fact that $Y \rightarrow Z$ is a generically 10-1 map is closely related to the fact that $\mathcal{E}$ is a projective variety of dimension 5 and degree 10. Although the individual linear equations $\mathbf{y}_i^\top \mathbf{E} \mathbf{x}_i = 0$ are special, considered together they determine a dominant rational map $Z \dashrightarrow \mathbb{G}_{3,8}$. We prove this in the next Proposition; a different proof may be found in [65, Sec. 5.2.3].

Proposition 4.1. The rational map

$$\begin{aligned} Z &\dashrightarrow \mathbb{G}_{3,8} \\ (\mathbf{x}_1, \dots, \mathbf{y}_5) &\mapsto \{\mathbf{E} \mid \mathbf{y}_i^\top \mathbf{E} \mathbf{x}_i = 0\} \end{aligned}$$

is dominant.

Proof. First observe that $\dim Z = \dim \mathbb{G}_{3,8} = 20$. To show that the map is dominant, it is enough to exhibit a point in the domain where the Jacobian of this map, expressed in local coordinates, is a 20×20 matrix of full rank. We verify this explicitly, using Stiefel coordinates on $\mathbb{G}_{3,8}$, with Macaulay2. $\square$

As a corollary, we deduce a geometric version of Hartley, Nistér, and Stewenius’s result [71, Section 5] that the Galois group of the five-point problem in its essential matrix formulation is the full symmetric group S_{10} .

Corollary 4.2. $\text{Mon}(Y/Z) \cong S_{10}$.

Proof. The branched cover

$$\{(\mathbf{E}, L) \in \mathcal{E} \times \mathbb{G}_{3,8} \mid \mathbf{E} \in L\} \rightarrow \mathbb{G}_{3,8}$$

has full-symmetric Galois/monodromy group S_{10} , as follows from a general fact sometimes known as the uniform position principle [5, pp.111–112]. The monodromy loops in $\mathbb{G}_{3,8}$ that generate S_{10} may be assumed to lie in a regular locus in the image of $Z \dashrightarrow \mathbb{G}_{3,8}$, and thus pull back to monodromy loops in Z which act as S_{10} on the fiber Y_z . $\square$

In formulating the five-point problem, it is possible to normalize the translations and depths in various ways. For instance, we might use the normalization $\|\mathbf{t}\|^2 = 1$, giving rise to a branched cover $W \rightarrow Z$ with $W \subset \mathrm{SO}_{\mathbb{C}}(3) \times \mathbb{C}^{13} \times Z$. More explicitly,

$$W = \{(\mathbf{R}, (\mathbf{t}, \alpha_1, \dots, \beta_5), (\mathbf{x}_1, \dots, \mathbf{y}_5)) \in \mathrm{SO}_{\mathbb{C}}(3) \times \mathbb{C}^{13} \times Z \mid (1) \text{ holds and } \|\mathbf{t}\|^2 = 1\}.$$

There is a map $W \rightarrow X$ induced by the quotient map which sends $(\mathbf{t}, \alpha_1, \dots, \beta_5)$ to its projective equivalence class in $\mathbb{P}^{12}$

$$\begin{aligned} p : \{(\mathbf{t}, \alpha_1, \dots, \beta_5) \mid \|\mathbf{t}\|^2 = 1\} &\rightarrow \mathbb{P}^{12} \\ (\mathbf{t}, \alpha_1, \dots, \beta_5) &\mapsto [(\mathbf{t}, \alpha_1, \dots, \beta_5)] \end{aligned}$$

The fiber in W over generic $x = (\mathbf{R}, [\mathbf{t}, \alpha_1, \dots, \beta_5], (\mathbf{x}_1, \dots, \mathbf{y}_5)) \in X$ is

$$W_x = \{(\mathbf{R}, (\mathbf{t}, \alpha_1, \dots, \beta_5), (\mathbf{x}_1, \dots, \mathbf{y}_5)), (\mathbf{R}, (-\mathbf{t}, -\alpha_1, \dots, -\beta_5), (\mathbf{x}_1, \dots, \mathbf{y}_5))\}.$$

Altogether, the branched covers of the five-point problem are related by the sequence of maps

$$W \rightarrow X \rightarrow Y \rightarrow Z.$$

Arguably, solutions in W carry more physical meaning than solutions in X , since a solution is physically meaningful only if all world points are in front of the camera, meaning $\alpha_1, \dots, \beta_5 > 0$.

When solving the five-point problem in practice, the essential matrix can only be determined up to scale. Each essential matrix solution corresponds to a point $y \in Y$, and the fiber X_y associates two consistent sets of rotations, translations, and depths. In X , translations and depths are regarded as elements of the projective space $\mathbb{P}^{12}$, and a chosen representative in homogeneous coordinates need not satisfy $\alpha_1, \dots, \beta_5 > 0$. In W , a solution $(\mathbf{R}, \mathbf{t}, \alpha_1, \dots, \beta_5)$ with all depths negative maps to another solution with all depths positive under the deck transformation

$$\begin{aligned} \Psi_{\mathrm{sgn}} : W &\rightarrow W \\ ((\mathbf{R}, \mathbf{t}, \alpha_1, \dots, \beta_5), (\mathbf{x}_1, \dots, \mathbf{y}_5)) &\mapsto ((\mathbf{R}, -\mathbf{t}, -\alpha_1, \dots, -\beta_5), (\mathbf{x}_1, \dots, \mathbf{y}_5)). \end{aligned}$$

The four possibilities for $(\mathbf{R}, \mathbf{t}, \alpha_1, \dots, \beta_5) \in \mathrm{SO}_{\mathbb{C}}(3) \times \mathbb{C}^{13}$ corresponding to a given essential matrix are illustrated in [38, Figure 9.12]. They correspond to the four deck transformations of $W \rightarrow Z$: the identity map $W \rightarrow W$, the sign-symmetry Ψ_{sgn} , a twisted pair Ψ_W , and the composition $\Psi_{\mathrm{sgn}} \circ \Psi_W$.

The four elements in the generic fiber of $W \rightarrow Y$ are illustrated in [38, Figure 9.12]. The various Galois/monodromy groups for the five-point problem are summarized in Result 4.2.

Result 4.2. The Galois/monodromy groups of the five-point problem are as follows:

$$\begin{aligned} \mathrm{Mon}(W/Z) &\cong (C_2)^9 \rtimes (S_2 \wr S_{10}) \hookrightarrow S_{40} \\ \mathrm{Mon}(X/Z) &\cong S_2 \wr S_{10} \cap A_{20} \hookrightarrow S_{20} \\ \mathrm{Mon}(Y/Z) &\cong S_{10}, \end{aligned}$$

where C_2^9 is the subgroup of $(\tau_1, \dots, \tau_{20}) \in (S_2)^{20}$ generated by $\tau_{2i-1}\tau_{2i}$ for $i = 1, \dots, 20$.

4.2 Two-view homography

We now move on to a minimal problem for planar scenes. We use similar notation as in the five-point problem, with constraints

$$\begin{aligned} \mathbf{R}^\top \mathbf{R} &= \mathbf{I}, \quad \det \mathbf{R} = 1, \\ \beta_i \mathbf{y}_i &= \mathbf{R} \alpha_i \mathbf{x}_i + \mathbf{t}, \quad \alpha_i, \beta_i \neq 0, \quad i = 1, \dots, 4, \\ \det \left(\begin{bmatrix} \alpha_1 \mathbf{x}_1 & \alpha_2 \mathbf{x}_2 & \alpha_3 \mathbf{x}_3 & \alpha_4 \mathbf{x}_4 \\ 1 & 1 & 1 & 1 \end{bmatrix} \right) &= 0. \end{aligned} \tag{20}$$

Our incidence variety X is given by

$$X \subseteq \mathrm{SO}_{\mathbb{C}}(3) \times \mathbb{P}_{\mathbb{C}}^{10} \times \underbrace{(\mathbb{C}^2 \times \{1\})^4 \times (\mathbb{C}^2 \times \{1\})^4}_Z$$

such that Equations 20 hold for all $(\mathbf{R}, (\mathbf{t}, \alpha_1, \dots, \alpha_4, \beta_1, \dots, \beta_4), (\mathbf{x}_1, \dots, \mathbf{x}_4, \mathbf{y}_1, \dots, \mathbf{y}_4)) \in X$. Projection of X onto Z defines a branched cover of degree 12. This branched cover is birationally equivalent to the joint camera map of the problem $\begin{smallmatrix} \bullet & \bullet \\ \vdots & \vdots \\ \bullet & \bullet \end{smallmatrix}$ from [20], since the fifth point on both lines in each image is generically determined from the other four points. Result 4.1 tells us the Galois/monodromy group is $C_2 \times C_2 \rtimes (S_2 \wr S_3 \cap A_6)$. The GAP command `MinimalGeneratingSet` shows that this group is minimally generated by two permutations: in cycle notation,

$$\mathrm{Mon}(X/Z) \cong \langle (1 \ 2)(3 \ 4)(5 \ 12 \ 8 \ 9)(6 \ 11 \ 7 \ 10), (1 \ 11 \ 5)(2 \ 10 \ 8)(3 \ 9 \ 7)(4 \ 12 \ 6) \rangle. \tag{21}$$




Figure 5: Correspondence between block systems (left) and intermediate fields (right) for the calibrated homography problem. The notation K_H means the intermediate field of an extension K/F fixed elementwise by a subgroup $H \leq \mathrm{Aut}(K/F)$.

The lattice of block systems is depicted on the left in Figure 5. The vertex labels correspond to stabilizer subgroups of $\mathrm{Mon}(X/Z)$, and the edges are labeled by the degrees of maps appearing in some decomposition of the form in Equation (12). To the right is the inverted lattice of intermediate fields. Like the majority of examples in this paper, $\mathbb{C}(X)/\mathbb{C}(Z)$ is not a Galois extension.

Before we determine a decomposition, we first describe the group of deck transformations. The centralizer in S_{12} is

$$\left\langle (1\ 3)(2\ 4)(5\ 7)(6\ 8)(9\ 11)(10\ 12), (1\ 4)(2\ 3)(5\ 8)(6\ 7)(9\ 12)(10\ 11) \right\rangle \cong C_2 \times C_2.$$

The deck transformation corresponding to the first generator is the twisted pair map Ψ_1 , defined just as in Equation (6). The second is a reflection-rotation symmetry Ψ_2 depicted in Figure 6.




Figure 6: Reflection-rotation symmetry for Equations 20

To get a formula for Ψ_2 , it is convenient to work with the equation of the unknown plane:

$$\langle \mathbf{n}, \mathbf{X} \rangle = d. \quad (22)$$

Note that $\mathbf{n}$ and d depend rationally on the data. The formula for Ψ_2 is given by

$$\begin{aligned} \Psi_2(\mathbf{R}) &= \mathbf{R} \left(2 \frac{\mathbf{n}\mathbf{n}^\top}{\mathbf{n}^\top \mathbf{n}} - \mathbf{I} \right) \\ \Psi_2(\mathbf{t}) &= -\mathbf{t} - \frac{2d}{\mathbf{n}^\top \mathbf{n}} \mathbf{R}\mathbf{n} \\ \Psi_2(\alpha_i) &= \alpha_i \\ \Psi_2(\beta_i) &= -\beta_i \end{aligned} \quad (23)$$

$$\Psi_2(\mathbf{x}_1, \dots, \mathbf{x}_4, \mathbf{y}_1, \dots, \mathbf{y}_4) = (\mathbf{x}_1, \dots, \mathbf{x}_4, \mathbf{y}_1, \dots, \mathbf{y}_4).$$

To better understand the effect of Ψ_2 on $\mathbf{t}$, let $\mathbf{X}$ be any point on the scene plane and calculate

$$-\mathbf{t} - \frac{2d}{\mathbf{n}^\top \mathbf{n}} \mathbf{R}\mathbf{n} = -\mathbf{t} - \mathbf{R}\mathbf{X} - \mathbf{R} \left(2 \frac{\mathbf{n}\mathbf{n}^\top}{\mathbf{n}^\top \mathbf{n}} - \mathbf{I} \right) \mathbf{X} \quad (24)$$

$$= -\mathbf{R} \left(2 \frac{\mathbf{n}\mathbf{n}^\top}{\mathbf{n}^\top \mathbf{n}} - \mathbf{I} \right) \left(\left(\mathbf{I} - 2 \frac{\mathbf{n}\mathbf{n}^\top}{\mathbf{n}^\top \mathbf{n}} \right) (-\mathbf{R}^\top \mathbf{t} - \mathbf{X}) + \mathbf{X} \right) \quad (25)$$

$$= -\Psi_2(\mathbf{R}) \left(\left(\mathbf{I} - 2 \frac{\mathbf{n}\mathbf{n}^\top}{\mathbf{n}^\top \mathbf{n}} \right) (-\mathbf{R}^\top \mathbf{t} - \mathbf{X}) + \mathbf{X} \right) \quad (26)$$

This may be understood as follows: we take $-\mathbf{R}^\top \mathbf{t}$, which is the center of the second camera expressed in the frame of the first camera (cf. Eq. 20), then reflect this vector through the plane and transform it back to the vector representing the center of the first camera expressed in the frame of the reflected second camera (by multiplying by $-\Psi_2(\mathbf{R})$).

Proposition 4.3. Ψ_1 and Ψ_2 generate the deck transformation group for the planar calibrated homography problem $X \rightarrow Z$ defined by Equations 20. The corresponding permutations which centralize $\text{Mon}(X/Z)$ are as follows:

$$\begin{aligned}\Psi_1 \circ \Psi_2 &\leftrightarrow (1\ 2)(3\ 4)(5\ 6)(7\ 8)(9\ 10)(11\ 12) \\ \Psi_1 &\leftrightarrow (1\ 3)(2\ 4)(5\ 7)(6\ 8)(9\ 11)(10\ 12) \\ \Psi_2 &\leftrightarrow (1\ 4)(2\ 3)(5\ 8)(6\ 7)(9\ 12)(10\ 11)\end{aligned}$$

These correspond to the maximal chains in the lattice of block systems (see Figure 5).

Proof. We verify that Ψ_1 and Ψ_2 really are deck transformations. The rest is elementary or follows from Proposition 2.16. Appealing to well-known properties of the twisted pair Ψ_1 , it suffices for us to check that planarity of the scene is preserved:

$$\det \begin{pmatrix} \Psi_1(\alpha_1)\mathbf{x}_1 & \Psi_1(\alpha_2)\mathbf{x}_2 & \Psi_1(\alpha_3)\mathbf{x}_3 & \Psi_1(\alpha_4)\mathbf{x}_4 \\ 1 & 1 & 1 & 1 \end{pmatrix} = 0.$$

Letting $\mathbf{m} = \frac{2}{\mathbf{t}^\top \mathbf{t}} \mathbf{R}^\top \mathbf{t}$, we may compute this determinant as follows:

$$\begin{aligned} \left(\prod_{i=1}^4 (1 + \mathbf{m}^\top \alpha_i \mathbf{x}_i) \right)^{-1} \det \begin{pmatrix} \alpha_1 \mathbf{x}_1 & \alpha_2 \mathbf{x}_2 & \alpha_3 \mathbf{x}_3 & \alpha_4 \mathbf{x}_4 \\ 1 + \mathbf{m}^\top \alpha_1 \mathbf{x}_1 & 1 + \mathbf{m}^\top \alpha_2 \mathbf{x}_2 & 1 + \mathbf{m}^\top \alpha_3 \mathbf{x}_3 & 1 + \mathbf{m}^\top \alpha_4 \mathbf{x}_4 \end{pmatrix} = \\ \left(\prod_{i=1}^4 (1 + \mathbf{m}^\top \alpha_i \mathbf{x}_i) \right)^{-1} \det \begin{pmatrix} \alpha_1 \mathbf{x}_1 & \alpha_2 \mathbf{x}_2 & \alpha_3 \mathbf{x}_3 & \alpha_4 \mathbf{x}_4 \\ 1 & 1 & 1 & 1 \end{pmatrix} = 0. \end{aligned}$$

For Ψ_2 , it is clear that planarity of the scene is preserved and that $\Psi_2(\mathbf{R}) \in \text{SO}_{\mathbb{C}}(3)$. If we substitute $\Psi_2(x)$ into the point correspondence constraint in (20) and take $\mathbf{X} = \alpha_i \mathbf{x}_i$ in Equation (24), then

$$-\beta_i \mathbf{y}_i = \mathbf{R} \left(2 \frac{\mathbf{n} \mathbf{n}^\top}{\mathbf{n}^\top \mathbf{n}} - \mathbf{I} \right) \alpha_i \mathbf{x}_i + \left(-\mathbf{t} - \mathbf{R} \alpha_i \mathbf{x}_i - \mathbf{R} \left(2 \frac{\mathbf{n} \mathbf{n}^\top}{\mathbf{n}^\top \mathbf{n}} - \mathbf{I} \right) \alpha_i \mathbf{x}_i \right) = -\mathbf{R} \alpha_i \mathbf{x}_i - \mathbf{t}.$$

We conclude that equations (20) are invariant up to sign under application of Ψ_2 . $\square$

Finally, we describe a decomposition of $X \rightarrow Z$

$$X \dashrightarrow Y_1 \dashrightarrow Y_2 \rightarrow Z, \quad (27)$$

corresponding to the left-most chain in Figure 5. This decomposition makes use of the *calibrated homography matrix* associated to $(\mathbf{R}, \mathbf{t})$ and the scene plane:

$$\mathbf{H} = \mathbf{R} + \frac{1}{d} \mathbf{t} \mathbf{n}^\top. \quad (28)$$

Up to scale, any 3×3 matrix has the form (28). On the other hand, any real calibrated homography matrix has an eigenvalue equal to 1 (see eg. [63, Lemma 5.18]), and thus lies on an irreducible hypersurface of degree 6:

$$\mathcal{H}_1 = \{\mathbf{H} \in \mathbb{C}^{3 \times 3} \mid \det(\mathbf{H}^\top \mathbf{H} - \mathbf{I}) = 0\}.$$

In our decomposition, we may take

$$Y_1 = \{(\mathbf{H}, ((\mathbf{x}_1, \dots, \mathbf{x}_4), (\mathbf{y}_1, \dots, \mathbf{y}_4))) \in \mathcal{H}_1 \times (\mathbb{P}^2)^4 \mid \mathbf{y}_i \sim \mathbf{H}\mathbf{x}_i, i = 1, \dots, 4\}.$$

Here we use the standard notation $\sim$ indicate that two vectors are equal up to scale. We note that each of these correspondence constraints is equivalent to the vanishing of three homogeneous, non-independent linear equations

$$[\mathbf{y}_i]_{\times} \mathbf{H}\mathbf{x}_i = 0.$$

Short calculations reveal that $x \in X_z$ and $\Psi_1 \circ \Psi_2(x) \in X_z$ map to the same point in Y_1 , and that the homography matrices associated to x and $\Psi_2(x)$ differ only up to sign. We also note that Y_1 is irreducible, since it is a Zariski-open subset of the graph of $\mathcal{H}_1 \times (\mathbb{P}^2)^4 \dashrightarrow (\mathbb{P}^2)^4$.

The projection $Y_1 \rightarrow Z$ has a deck transformation given by the sign-symmetry $\mathbf{H} \mapsto -\mathbf{H}$. To remove this last symmetry, we define

$$\begin{aligned} s &= \frac{1}{\mathbf{H}_{1,1}^2} \\ \mathbf{S} &= \frac{1}{\mathbf{H}_{1,1}} \mathbf{H} \end{aligned} \quad (29)$$

and take $Y_1 \dashrightarrow Y_2 \subset \mathbb{C}^9 \times (\mathbb{P}^2)^4 \times (\mathbb{P}^2)^4$ by mapping $\mathbf{H}$ to s and the 8 non-constant entries of $\mathbf{S}$. For generic data $(\mathbf{x}_1, \dots, \mathbf{y}_4) \in Z$, the ideal in $\mathbb{C}[s, \mathbf{S}_{1,2}, \mathbf{S}_{1,3}, \mathbf{S}_{2,1}, \mathbf{S}_{2,2}, \mathbf{S}_{2,3}, \mathbf{S}_{3,1}, \mathbf{S}_{3,2}, \mathbf{S}_{3,3}]$ given by

$$\langle \det(\mathbf{S}^\top \mathbf{S} - s\mathbf{I}), [\mathbf{y}_1]_{\times} \mathbf{S}\mathbf{x}_1, [\mathbf{y}_2]_{\times} \mathbf{S}\mathbf{x}_2, [\mathbf{y}_3]_{\times} \mathbf{S}\mathbf{x}_3, [\mathbf{y}_4]_{\times} \mathbf{S}\mathbf{x}_4 \rangle \quad (30)$$

has dimension 0 and degree $3 = \deg(Y_2/Z)$. The algebraic complexity as captured by the Galois group matches that of a well-known algorithm for computing $\mathbf{H}$, in which one must compute the singular values of a 3×3 matrix $\lambda\mathbf{H}$ recovered up to scale from the four point correspondences (see eg. [38, Algorithm 4.1] or [63, Algorithm 5.2].)

4.3 Three-view homography

Finally, we consider the minimal problem $\boxed{\bullet, \bullet, \bullet}$, where the task is to recover the relative orientation of three cameras from the input data of four point correspondences which lie on the incidence variety

$$Z = \{(\mathbf{x}_1, \dots, \mathbf{x}_4, \mathbf{y}_1, \dots, \mathbf{y}_4, \mathbf{z}_1, \dots, \mathbf{z}_4) \in (\mathbb{P}^2)^{12} \mid \mathbf{x}_3 \in \overline{\mathbf{x}_1 \mathbf{x}_2}, \mathbf{y}_3 \in \overline{\mathbf{y}_1 \mathbf{y}_2}, \mathbf{z}_3 \in \overline{\mathbf{z}_1 \mathbf{z}_2}\}.$$

Unlike in the previous section, there is no longer a twisted pair symmetry. However, there are two symmetries analogous to the deck transformation Ψ_2 . For this problem, the joint camera map defined in [20] is birationally equivalent to a branched cover whose fibers are pairs of homography matrices, which are compatible in the sense that they share the same normal vector. Thus, the solutions of interest lie on the subvariety $\mathcal{H}_2 \subset (\mathbb{C}^{3 \times 3})^2$ defined to be the closed image of the map

$$\begin{aligned} (\mathrm{SO}_{\mathbb{C}}(3))^2 \times (\mathbb{C}^3)^3 &\rightarrow (\mathbb{C}^{3 \times 3})^2 \\ (\mathbf{R}_1, \mathbf{R}_2, \mathbf{t}_1, \mathbf{t}_2, \mathbf{n}) &\mapsto (\mathbf{R}_1 + \mathbf{t}_1 \mathbf{n}^\top, \mathbf{R}_2 + \mathbf{t}_2 \mathbf{n}^\top). \end{aligned}$$

Notice that, unlike in (28), we have absorbed the constant $\frac{1}{d}$ into $\mathbf{t}$ for each homography matrix. We wish to compute the fibers of the branched cover $X \rightarrow Z$, where

$$X = \{((\mathbf{H}_1, \mathbf{H}_2), (\mathbf{x}_1, \dots, \mathbf{z}_4)) \in \mathcal{H}_2 \times Z \mid \mathbf{H}_1 \mathbf{x}_i \sim \mathbf{y}_i, \mathbf{H}_2 \mathbf{x}_i \sim \mathbf{z}_i, i = 1, \dots, 4\}. \quad (31)$$

We have $\deg(X/Z) = 64$, and Result 4.1 tells us $\text{Mon}(X/Z) \cong S_2 \wr (S_2 \wr S_{16} \cap A_{32}) \cap A_{64}$. It follows that there exists a decomposition

$$X \dashrightarrow Y_1 \dashrightarrow Y_2 \dashrightarrow Z$$

with $\deg(X/Y_1) = \deg(Y_1/Y_2) = 2$ and $\deg(Y_2/Z) = 16$. The deck transformations of $X \rightarrow Z$ are easily seen to be $(\mathbf{H}_1, \mathbf{H}_2) \mapsto (\pm \mathbf{H}_1, \pm \mathbf{H}_2)$. Thus, we may use separating invariants for this linear group action as in the previous section to write down the maps $X \dashrightarrow Y_1$ and $Y_1 \dashrightarrow Y_2$.

However, our description of X is unsatisfying from the point of view of constructing polynomial solvers, since we have only described $\mathcal{H}_2$ parametrically. We leave determining the ideal $\mathcal{I}_{\mathcal{H}_2}$ as a challenging open problem in algebraic vision, analogous to previous works [2, 3]. Our final Result 4.3 is a partial solution to this implicitization problem, which describes an ideal contained in $\mathcal{I}_{\mathcal{H}_2}$. The generators of this ideal and the linear correspondence constraints in (31) generate a 0-dimensional ideal of degree 64 for generic data $z = (\mathbf{x}_i, \mathbf{y}_i, \mathbf{z}_i)$.

Drawing on the description of $\mathcal{H}_1$ from the previous section, consider the map

$$\begin{aligned} \mathcal{H}_2 &\rightarrow (\mathbb{C}^{3 \times 3})^2 \\ (\mathbf{H}_1, \mathbf{H}_2) &\mapsto (\mathbf{H}_1^\top \mathbf{H}_1 - \mathbf{I}, \mathbf{H}_2^\top \mathbf{H}_2 - \mathbf{I}). \end{aligned}$$

Writing $\mathbf{W}_i = \mathbf{H}_i^\top \mathbf{H}_i - \mathbf{I}$ for $i = 1, 2$, we compute the vanishing ideal of the image of this map in $\mathbb{C}[\mathbf{W}_1, \mathbf{W}_2]$. This ideal is generated by 4 cubics and 15 quartics. The cubic constraints obtained are

$$\det(\mathbf{W}_1) = \det(\mathbf{W}_2) = \det(\mathbf{W}_1 + \mathbf{W}_2) = \det(\mathbf{W}_1 - \mathbf{W}_2) = 0. \quad (32)$$

In what follows, it is enough for us to consider two of the 15 quartics in $\mathbb{C}[\mathbf{W}_1, \mathbf{W}_2]$, which have alternate expressions in terms of resultants which enforce the compatibility between homographies:

$$\begin{aligned} \text{Res}_{n_1}(\mathbf{W}_{1,3,3} n_1^2 - 2\mathbf{W}_{1,1,3} n_1 + \mathbf{W}_{1,1,1}, \mathbf{W}_{2,3,3} n_1^2 - 2\mathbf{W}_{2,1,3} n_1 + \mathbf{W}_{2,1,1}) &= 0, \\ \text{Res}_{n_2}(\mathbf{W}_{1,3,3} n_2^2 - 2\mathbf{W}_{1,2,3} n_2 + \mathbf{W}_{1,2,2}, \mathbf{W}_{2,3,3} n_2^2 - 2\mathbf{W}_{2,2,3} n_2 + \mathbf{W}_{2,2,2}) &= 0, \end{aligned} \quad (33)$$

where $\mathbf{n} = (n_1, n_2, n_3)$.

Substituting $\mathbf{W}_i = \mathbf{H}_i^\top \mathbf{H}_i - \mathbf{I}$ into (32) yields four polynomials of degree 6 vanishing on $\mathcal{H}_2$. Using Bertini [11], we computed points where these equations vanish by tracking $6^4 = 1296$ homotopy paths. Out of these points, 336 lie on $\mathcal{H}_2$. The remaining 960 paths resulted in 224 points not on $\mathcal{H}_2$, each occurring with multiplicity 4. We confirmed that the degree of the variety $\mathcal{H}_2$ is indeed 336 using monodromy.

Unlike in the five-point problem, the linear equations implied by $\mathbf{H}_1 \mathbf{x}_i \sim \mathbf{y}_i$, $\mathbf{H}_2 \mathbf{x}_i \sim \mathbf{z}_i$ are non-generic. The number of solutions to these linear equations and the four degree-6 equations obtained from (32) is 320. Out of these solutions, only 84 satisfy the degree-8 equations obtained from (33). To obtain the degree 64 reported in [20], it is sufficient to impose the additional constraint $\det \mathbf{H}_1 \neq 0$. In summary, we have the following result.

Result 4.3. For each $z = (\mathbf{x}_1, \dots, \mathbf{z}_4) \in Z$, let $I_z \subset \mathbb{C}[\mathbf{H}_1, \mathbf{H}_2, D]$ be the ideal in 19 variables generated by the linear relations $\mathbf{H}_1 \mathbf{x}_i \sim \mathbf{y}_i$, $\mathbf{H}_2 \mathbf{x}_i \sim \mathbf{z}_i$ for $i = 1, \dots, 4$, the saturation constraint $D \det \mathbf{H}_1 - 1 = 0$, and six equations obtained by setting $(\mathbf{W}_1, \mathbf{W}_2) = (\mathbf{H}_1^\top \mathbf{H}_1 - \mathbf{I}, \mathbf{H}_2^\top \mathbf{H}_2 - \mathbf{I})$ in equations (32) and (33). For generic $z \in Z$, we have $\deg(I_z) = 64$. Moreover, after substituting $s_1, s_2, \mathbf{S}_1, \mathbf{S}_2$ as in (29), (30), we obtain an ideal of degree 16 for generic data.

5 Conclusion and outlook

Galois/monodromy groups reveal the intrinsic algebraic structure of problems in enumerative geometry. We have shown how this structure can be revealed on both new and old examples from computer vision. We believe that numerical algebraic geometry and computational group theory are valuable additions to the arsenal of techniques used by researchers interested in building minimal solvers. Although these methods allow us to identify decomposable branched covers, the problem of automatically computing a decomposition as in Equation (12) seems hard in general. We have shown how understanding the underlying geometry allows us to solve it in cases of interest.

One potential extension would be to study symmetries of non-minimal problems. Here, one could consider branched covers $X \dashrightarrow Z$ where the base space Z is a proper subvariety of some affine space. For instance, the base space for calibrated relative pose from six point correspondences would be the Chow variety computed in [29]. We still expect a finite field extension $\mathbb{C}(X)/\mathbb{C}(Z)$ when there are more than the minimal number of observations. The main difficulty here is how to construct monodromy loops in Z , since a straight-line homotopy will no longer work. One possible remedy would be to project Z onto an affine space and apply Part 2) of Proposition 2.11.

Proposition 2.16 implies an effective criterion for detecting symmetries of minimal problems; in this paper, we determined formulas for these symmetries in an *ad-hoc* fashion. However, this result also suggests the possibility of systematically computing these deck transformations by rational function interpolation together from the data obtained from monodromy computations. We plan to report on such a systematic approach in upcoming work.

Acknowledgements

We are grateful to ICERM (NSF DMS-1439786 and the Simons Foundation grant 507536) for their support and for hosting three events where significant progress on this project was made: the Fall 2018 semester program on Nonlinear Algebra, the Winter 2019 Algebraic Vision research cluster, and the Fall 2020 virtual workshop on Monodromy and Galois groups in enumerative geometry and applications. V. Korotynskiy and T. Pajdla were supported by the EU Structural and Investment Funds, Operational Programme Research, Development and Education under the project IMPACT (reg. no. CZ.02.1.01/0.0/0.0/15_003/0000468), Grant Agency of the CTU in Prague project SGS21/056/OHK3/1T/13, and the EU H2020 SPRING no. 871245 projects. T. Duff acknowledges support from an NSF Mathematical Sciences Postdoctoral Research Fellowship (DMS-2103310), as well as partial support from NSF DMS #1719968 and #2001267. M. Regan was supported in part by Schmitt Leadership Fellowship in Science and Engineering and NSF grant CCF-1812746. We are particularly grateful to Jon Hauenstein, Anton Leykin, and Frank Sottile for helpful suggestions.

References

- [1] Sameer Agarwal, Hon-leung Lee, Bernd Sturmfels, and Rekha R. Thomas, *On the existence of epipolar matrices*, International Journal of Computer Vision **121** (2017), no. 3, 403–415.
- [2] Sameer Agarwal, Andrew Pryhuber, and Rekha R. Thomas, *Ideals of the multiview variety*, IEEE Trans. Pattern Anal. Mach. Intell. **43** (2021), no. 4, 1279–1292.

- [3] Chris Aholt and Luke Oeding, *The ideal of the trifocal variety*, Math. Comp. **83** (2014), no. 289, 2553–2574. MR 3223346
- [4] Andre Ameller, Bill Triggs, and Long Quan, *Camera pose revisited: New linear algorithms*, 2002.
- [5] E. Arbarello, M. Cornalba, P. A. Griffiths, and J. Harris, *Geometry of algebraic curves. Vol. I*, Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], vol. 267, Springer-Verlag, New York, 1985. MR 770932
- [6] Erik Ask, Yubin Kuang, and Kalle Åström, *Exploiting p -fold symmetries for faster polynomial equation solving*, Proceedings of the 21st International Conference on Pattern Recognition, ICPR 2012, Tsukuba, Japan, November 11-15, 2012, IEEE Computer Society, 2012, pp. 3232–3235.
- [7] Chad Awtrey, Nakhila Mistry, and Nicole Soltz, *Centralizers of transitive permutation groups and applications to Galois theory*, Missouri J. Math. Sci. **27** (2015), no. 1, 16–32. MR 3431112
- [8] Daniel Barath, *Five-point fundamental matrix estimation for uncalibrated cameras*, 2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018, 2018, pp. 235–243.
- [9] Daniel Barath and Levente Hajder, *Efficient recovery of essential matrix from two affine correspondences*, IEEE Trans. Image Processing **27** (2018), no. 11, 5328–5337.
- [10] Daniel Barath, Tekla Toth, and Levente Hajder, *A minimal solution for two-view focal-length estimation using two affine correspondences*, 2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017, 2017, pp. 2557–2565.
- [11] Daniel J. Bates, Jonathan D. Hauenstein, Andrew J. Sommese, and Charles W. Wampler, *Bertini: Software for Numerical Algebraic Geometry*, Available at bertini.nd.edu with permanent doi: [dx.doi.org/10.7274/R0H41PB5](https://doi.org/10.7274/R0H41PB5).
- [12] ———, *Numerically solving polynomial systems with Bertini*, SIAM, 2013.
- [13] Paul Breiding and Sascha Timme, *HomotopyContinuation.jl: A Package for Homotopy Continuation in Julia*, Mathematical Software – ICMS 2018 (Cham), Springer International Publishing, 2018, pp. 458–465.
- [14] Taylor Brysiewicz, Jose Israel Rodriguez, Frank Sottile, and Thomas Yahl, *Solving decomposable sparse systems*, Numerical Algorithms (2021).
- [15] Martin Byröd, Klas Josephson, and Kalle Åström, *A column-pivoting based strategy for monomial ordering in numerical Gröbner basis calculations*, European Conference on Computer Vision (ECCV), vol. 5305, Springer, 2008, pp. 130–143.
- [16] F. Camposeco, T. Sattler, and M. Pollefeys, *Minimal solvers for generalized pose and scale estimation from two rays and one point*, ECCV – European Conference on Computer Vision, 2016, pp. 202–218.
- [17] Robert M. Corless, Karin Gatermann, and Ilias S. Kotsireas, *Using symmetries in the eigenvalue method for polynomial systems*, J. Symbolic Comput. **44** (2009), no. 11, 1536–1550. MR 2561287

- [18] Fernando Cukierman, *Monodromy of projections*, Mat. Contemp. **16** (1999), 9–30, 15th School of Algebra (Portuguese) (Canela, 1998). MR 1756825
- [19] Michel Demazure, *Sur deux problemes de reconstruction*, Tech. Report RR-0882, INRIA, July 1988, url: <https://hal.inria.fr/inria-00075672>.
- [20] T. Duff, K. Kohn, A. Leykin, and T. Pajdla, *PLMP - point-line minimal problems in complete multi-view visibility*, 2019 IEEE/CVF International Conference on Computer Vision (ICCV), 2019, pp. 1675–1684.
- [21] Timothy Duff, Cvetelina Hill, Anders Jensen, Kisun Lee, Anton Leykin, and Jeff Sommars, *Solving polynomial systems via homotopy continuation and monodromy*, IMA J. Numer. Anal. **39** (2019), no. 3, 1421–1446. MR 3984062
- [22] Timothy Duff, Cvetelina Hill, Anders Nedergaard Jensen, Kisun Lee, Anton Leykin, and Jeff Sommars, *MonodromySolver: solving polynomial systems via monodromy. Version 1.14*, A Macaulay2 package available at "<http://www.math.gatech.edu/leykin>".
- [23] Timothy Duff, Kathlén Kohn, Anton Leykin, and Tomás Pajdla, *PL₁P - point-line minimal problems under partial visibility in three views*, Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI (Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, eds.), Lecture Notes in Computer Science, vol. 12371, Springer, 2020, pp. 175–192.
- [24] Ali Elqursh and Ahmed M. Elgammal, *Line-based relative pose estimation*, The 24th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, Colorado Springs, CO, USA, 20–25 June 2011, IEEE Computer Society, 2011, pp. 3049–3056.
- [25] Alexander Esterov and Lionel Lang, *Sparse polynomial equations and other enumerative problems whose Galois groups are wreath products*, arXiv preprint arXiv:1812.07912 (2018).
- [26] Ricardo Fabbri, Timothy Duff, Hongyi Fan, Margaret H. Regan, David da Costa de Pinho, Elias P. Tsigaridas, Charles W. Wampler, Jonathan D. Hauenstein, Peter J. Giblin, Benjamin B. Kimia, Anton Leykin, and Tomás Pajdla, *TRPLP - Trifocal relative pose from lines at points*, 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020, IEEE, 2020, pp. 12070–12080.
- [27] Jean-Charles Faugere, Patrizia Gianni, Daniel Lazard, and Teo Mora, *Efficient computation of zero-dimensional Gröbner bases by change of ordering*, Journal of Symbolic Computation **16** (1993), no. 4, 329–344.
- [28] 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** (1981), no. 6, 381–395.
- [29] Gunnar Fløystad, Joe Kileel, and Giorgio Ottaviani, *The Chow form of the essential variety in computer vision*, J. Symbolic Comput. **86** (2018), 97–119. MR 3725215
- [30] The GAP Group, *GAP – Groups, Algorithms, and Programming, Version 4.11.1*, 2021.

- [31] Karin Gatermann, *Symbolic solution polynomial equation systems with symmetry*, Proceedings of the International Symposium on Symbolic and Algebraic Computation (New York, NY, USA), ISSAC '90, Association for Computing Machinery, 1990, p. 112–119.
- [32] Daniel R. Grayson and Michael E. Stillman, *Macaulay2, a software system for research in algebraic geometry*, Available at <http://www.math.uiuc.edu/Macaulay2/>.
- [33] J. A. Grunert, *Das pothenotische problem in erweiterter gestalt nebst über seine anwendungen in der geodäsie*, Grunerts Archiv für Mathematik und Physik **1** (1841), 238–248.
- [34] Jaime Gutierrez and David Sevilla, *Building counterexamples to generalizations for rational functions of Ritt's decomposition theorem*, J. Algebra **303** (2006), no. 2, 655–667. MR 2255128
- [35] Robert M. Haralick, Chung-Nan Lee, Karsten Ottenberg, and Michael Nölle, *Review and analysis of solutions of the three point perspective pose estimation problem*, Int. J. Comput. Vis. **13** (1994), no. 3, 331–356.
- [36] Joe Harris, *Galois groups of enumerative problems*, Duke Math. J. **46** (1979), no. 4, 685–724. MR 552521
- [37] R. Hartley and Hongdong Li, *An efficient hidden variable approach to minimal-case camera motion estimation*, IEEE PAMI **34** (2012), no. 12, 2303–2314.
- [38] Richard Hartley and Andrew Zisserman, *Multiple view geometry in computer vision*, 2nd ed., Cambridge, 2003.
- [39] Richard I. Hartley, *In defense of the eight-point algorithm*, IEEE Trans. Pattern Anal. Mach. Intell. **19** (1997), no. 6, 580–593.
- [40] Allen Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002. MR 1867354
- [41] Jonathan D Hauenstein and Margaret H Regan, *Adaptive strategies for solving parameterized systems using homotopy continuation*, Applied Mathematics and Computation **332** (2018), 19–34.
- [42] Jonathan D. Hauenstein and Margaret H. Regan, *Real monodromy action*, Applied Mathematics and Computation **373** (2020), 124983.
- [43] Jonathan D. Hauenstein, Jose Israel Rodriguez, and Frank Sottile, *Numerical computation of Galois groups*, Found. Comput. Math. **18** (2018), no. 4, 867–890. MR 3833644
- [44] Derek F. Holt, Bettina Eick, and Eamonn A. O'Brien, *Handbook of computational group theory*, Discrete Mathematics and its Applications (Boca Raton), Chapman & Hall/CRC, Boca Raton, FL, 2005. MR 2129747
- [45] Lars Hörmander, *An introduction to complex analysis in several variables*, revised ed., North-Holland Mathematical Library, Vol. 7, North-Holland Publishing Co., Amsterdam-London; American Elsevier Publishing Co., Inc., New York, 1973. MR 0344507
- [46] M. Husty and H.-P. Schröcker, *Algebraic geometry and kinematics*, Nonlinear Computational Geometry, volume 151 of The IMA Volumes in Mathematics and its Applications, Springer, 2010, pp. 85–107.

- [47] Gregor Kemper, *Separating invariants*, J. Symbolic Comput. **44** (2009), no. 9, 1212–1222. MR 2532166
- [48] Joe Kileel, *Minimal problems for the calibrated trifocal variety*, SIAM Journal on Applied Algebra and Geometry **1** (2017), no. 1, 575–598.
- [49] ———, *Minimal problems for the calibrated trifocal variety*, SIAM J. Appl. Algebra Geom. **1** (2017), no. 1, 575–598. MR 3705779
- [50] L. Kneip, R. Siegwart, and M. Pollefeys, *Finding the exact rotation between two images independently of the translation*, ECCV – European Conference on Computer Vision, 2012, pp. 696–709.
- [51] Erwin Kruppa, *Zur ermittlung eines objektes aus zwei perspektiven mit innerer orientierung*, Hölder, 1913.
- [52] Yubin Kuang and Kalle Åström, *Pose estimation with unknown focal length using points, directions and lines*, IEEE International Conference on Computer Vision, ICCV 2013, Sydney, Australia, December 1-8, 2013, 2013, pp. 529–536.
- [53] Yubin Kuang and Kalle Åström, *Stratified sensor network self-calibration from TDOA measurements*, 21st European Signal Processing Conference, 2013.
- [54] Zuzana Kukelova, Martin Bujnak, and Tomas Pajdla, *Automatic generator of minimal problem solvers*, European Conference on Computer Vision (ECCV), 2008.
- [55] Viktor Larsson and Kalle Åström, *Uncovering symmetries in polynomial systems*, Computer Vision – ECCV 2016 (Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, eds.), Springer International Publishing, 2016, pp. 252–267.
- [56] Viktor Larsson, Kalle Åström, and Magnus Oskarsson, *Efficient solvers for minimal problems by syzygy-based reduction*, Computer Vision and Pattern Recognition (CVPR), 2017.
- [57] Viktor Larsson, Kalle Åström, and Magnus Oskarsson, *Polynomial solvers for saturated ideals*, IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017, 2017, pp. 2307–2316.
- [58] Viktor Larsson, Zuzana Kukelova, and Yinqiang Zheng, *Making minimal solvers for absolute pose estimation compact and robust*, International Conference on Computer Vision (ICCV), 2017.
- [59] Vincent Lepetit, Francesc Moreno-Noguer, and Pascal Fua, *Epnnp: An accurate $O(n)$ solution to the pnp problem*, Int. J. Comput. Vis. **81** (2009), no. 2, 155–166.
- [60] Anton Leykin, *Homotopy continuation in Macaulay2*, International Congress on Mathematical Software, Springer, 2018, pp. 328–334.
- [61] Anton Leykin and Frank Sottile, *Galois groups of Schubert problems via homotopy computation*, Mathematics of Computation **78** (2009), 1749–1765.
- [62] Max Lieblich and Lucas Van Meter, *Two Hilbert schemes in computer vision*, SIAM Journal on Applied Algebra and Geometry **4** (2020), no. 2, 297–321.

- [63] Yi Ma, Stefano Soatto, Jana Kosecka, and S Shankar Sastry, *An invitation to 3-D vision: from images to geometric models*, vol. 26, Springer Science & Business Media, 2012.
- [64] Abraham Martín del Campo and Jose Israel Rodriguez, *Critical points via monodromy and local methods*, J. Symbolic Comput. **79** (2017), no. part 3, 559–574. MR 3563098
- [65] Stephen Maybank, *Theory of reconstruction from image motion*, vol. 28, Springer Science & Business Media, 2012.
- [66] Stephen J. Maybank and Olivier D. Faugeras, *A theory of self-calibration of a moving camera*, Int. J. Comput. Vis. **8** (1992), no. 2, 123–151.
- [67] Pedro Miraldo, Tiago Dias, and Srikumar Ramalingam, *A minimal closed-form solution for multi-perspective pose estimation using points and lines*, Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XVI, 2018, pp. 490–507.
- [68] Rick Miranda, *Algebraic curves and Riemann surfaces*, Graduate Studies in Mathematics, vol. 5, American Mathematical Society, Providence, RI, 1995. MR 1326604
- [69] Faraz M Mirzaei and Stergios I Roumeliotis, *Optimal estimation of vanishing points in a manhattan world*, International Conference on Computer Vision (ICCV), 2011.
- [70] David Mumford, *Algebraic geometry I: complex projective varieties*, Springer Science & Business Media, 1995.
- [71] D. Nister, R. Hartley, and H. Stewénus, *Using galois theory to prove structure from motion algorithms are optimal*, 2007 IEEE Conference on Computer Vision and Pattern Recognition, 2007, pp. 1–8.
- [72] David Nistér, *An efficient solution to the five-point relative pose problem*, IEEE transactions on pattern analysis and machine intelligence **26** (2004), no. 6, 756–770.
- [73] David Nistér and Frederik Schaffalitzky, *Four points in two or three calibrated views: Theory and practice*, Int. J. Comput. Vis. **67** (2006), no. 2, 211–231.
- [74] Dylan Peifer and Mahrud Sayrafi, *FGLM: Groebner bases via the FGLM algorithm. Version 1.1.0*, A Macaulay2 package available at <https://github.com/Macaulay2/M2/tree/master/M2/Macaulay2/packages>.
- [75] Long Quan and Zhong-Dan Lan, *Linear n-point camera pose determination*, IEEE Trans. Pattern Anal. Mach. Intell. **21** (1999), no. 8, 774–780.
- [76] R. Raguram, O. Chum, M. Pollefeys, J. Matas, and J.-M. Frahm, *USAC: A universal framework for random sample consensus*, IEEE Transactions on Pattern Analysis Machine Intelligence **35** (2013), no. 8, 2022–2038.
- [77] S. Ramalingam and P. F. Sturm, *Minimal solutions for generic imaging models*, CVPR – IEEE Conference on Computer Vision and Pattern Recognition, 2008.
- [78] Srikumar Ramalingam, Sofien Bouaziz, and Peter Sturm, *Pose estimation using both points and lines for geo-localization*, 2011 IEEE International Conference on Robotics and Automation, IEEE, 2011, pp. 4716–4723.

- [79] Gregory J. Reid, Jianliang Tang, and Lihong Zhi, *A complete symbolic-numeric linear method for camera pose determination*, Symbolic and Algebraic Computation, International Symposium ISSAC 2003, Drexel University, Philadelphia, Pennsylvania, USA, August 3-6, 2003, Proceedings (J. Rafael Sendra, ed.), ACM, 2003, pp. 215–223.
- [80] J. F. Ritt, *Prime and composite polynomials*, Trans. Amer. Math. Soc. **23** (1922), no. 1, 51–66. MR 1501189
- [81] Joseph J. Rotman, *An introduction to the theory of groups*, fourth ed., Graduate Texts in Mathematics, vol. 148, Springer-Verlag, New York, 1995. MR 1307623
- [82] Yohann Salaün, Renaud Marlet, and Pascal Monasse, *Robust and accurate line- and/or point-based pose estimation without manhattan assumptions*, European Conference on Computer Vision (ECCV), 2016.
- [83] Olivier Saurer, Marc Pollefeys, and Gim Hee Lee, *A minimal solution to the rolling shutter pose estimation problem*, Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on, IEEE, 2015, pp. 1328–1334.
- [84] Johannes Lutz Schönberger and Jan-Michael Frahm, *Structure-from-motion revisited*, Conference on Computer Vision and Pattern Recognition (CVPR), 2016.
- [85] Jean-Pierre Serre, *Topics in Galois theory*, second ed., Research Notes in Mathematics, vol. 1, A K Peters, Ltd., Wellesley, MA, 2008, With notes by Henri Darmon. MR 2363329
- [86] Igor R. Shafarevich, *Basic algebraic geometry. 1*, second ed., Springer-Verlag, Berlin, 1994, Varieties in projective space, Translated from the 1988 Russian edition and with notes by Miles Reid. MR 1328833
- [87] Noah Snavely, Steven M Seitz, and Richard Szeliski, *Modeling the world from internet photo collections*, International Journal of Computer Vision (IJCV) **80** (2008), no. 2, 189–210.
- [88] Andrew J. Sommese, Jan Verschelde, and Charles W. Wampler, *Numerical decomposition of the solution sets of polynomial systems into irreducible components*, SIAM J. Numer. Anal. **38** (2001), no. 6, 2022–2046. MR 1856241
- [89] Andrew J. Sommese and Charles W. Wampler, II, *The numerical solution of systems of polynomials*, World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, 2005, Arising in engineering and science. MR 2160078
- [90] H. Stewenius, C. Engels, and D. Nistér, *Recent developments on direct relative orientation*, ISPRS J. of Photogrammetry and Remote Sensing **60** (2006), 284–294.
- [91] Bill Triggs, *Camera pose and calibration from 4 or 5 known 3D points*, Proceedings of the International Conference on Computer Vision, Kerkyra, Corfu, Greece, September 20-25, 1999, IEEE Computer Society, 1999, pp. 278–284.
- [92] Lucas Van Meter, *A functorial approach to algebraic vision*, Ph.D. thesis, University of Washington, 2019.
- [93] Jonathan Ventura, Clemens Arth, and Vincent Lepetit, *An efficient minimal solution for multi-camera motion*, International Conference on Computer Vision (ICCV), 2015, pp. 747–755.

- [94] Yihong Wu and Zhanyi Hu, *PnP problem revisited*, J. Math. Imaging Vis. **24** (2006), no. 1, 131–141.
- [95] Henryk Żołądek, *The monodromy group*, Instytut Matematyczny Polskiej Akademii Nauk. Monografie Matematyczne (New Series) [Mathematics Institute of the Polish Academy of Sciences. Mathematical Monographs (New Series)], vol. 67, Birkhäuser Verlag, Basel, 2006. MR 2216496