Proving Optimality of Structure and Motion Algorithms Using Galois Theory

Richard Hartley
Australian National University
National ICT Australia
Canberra, Australia

David Nistér
Microsoft Live Labs
Microsoft Research
Seattle, WA, USA

Henrik Stewénius
Center for Visualization
Dep. of Computer Science
University of Kentucky, USA

April 27, 2008

Abstract

This paper presents a general method, based on Galois Theory, for establishing that a problem can not be solved by a 'machine' that is capable of the standard arithmetic operations, extraction of radicals (that is, m-th roots for any m) and Singular Value Decomposition, as well as extraction of roots of polynomials of degree smaller than n, but no other numerical operations.

The method is applied to two well known structure from motion problems: five point calibrated relative orientation, which can be realized by solving a tenth degree polynomial [9, 8], and L_2-optimal two-view triangulation, which can be realized by solving a sixth degree polynomial [5]. It is shown that both these solutions are optimal in the sense that an exact solution intrinsically requires the solution of a polynomial of the given degree (10 or 6 respectively), and cannot be solved by extracting roots of polynomials of any lesser degree.

1 Introduction

Many structure and motion problems can be reduced to the solution of a system of polynomial equations, and such systems of equations can in principle be reduced by elimination [3] to a single polynomial in one variable. If the polynomial is of degree 4 or less, then it may be solved in closed form by radicals (extraction of roots). For higher degree polynomials, numerical methods must generally be used. The number of solutions and an actual algorithm for solving the problem can be obtained by computing a Gröbner basis for the polynomial

equations [11, 12, 13]. This in principle gives a method for solving the problems, and as the papers just cited demonstrate, for many cases it gives excellent algorithms.

However, if the system of polynomials is large, this method can be complex and unstable, and solution of a polynomial of high degree is difficult in general. It is therefore advantageous to discover methods of solving structure and motion problems that require the solution of polynomials of as small degree as possible. The ideal is a solution that requires only the solution of 4-th degree equations, since the problem can then be solved in closed form (by radicals).

The purpose of the present paper is to present a method for placing a lower bound on the degree of such a polynomial needed to solve a particular problem. As examples of the technique, we consider the 5-point relative motion, and two-view triangulation problems. These can be solved by finding the roots of single polynomials of degree 10 or 6 respectively ([9, 5]). We show here that to solve the 5-point relative motion problem exactly, we essentially need to solve a polynomial of degree 10. No solution exists that involves only the solution of polynomials of lesser degree, even if we allow extraction of radicals (m-th roots) of any order. Similarly, the two-view L_2 triangulation problem requires the solution of a polynomial of degree 6. Since known algorithms exist involving solutions of polynomials of these degrees, these algorithms are optimal in the sense of the degree of the polynomial that needs to be solved. In particular, it follows that these problems have no solution in closed form by radicals.

The Singular Value Decomposition (SVD) is a popular and useful technique in structure from motion. Because of its reliability, algorithms that use SVD are often referred to as linear, although this is not strictly speaking a linear technique. We also show how the results of this paper can be extended to apply to SVD. Even adding this technique to our list of allowable operations can not avoid the necessity of solving polynomials of the indicated degree.

1.1 Brief Overview of the Proof

We briefly give an overview of what is to be proved, and the proof methodology. Recall that our goal is to demonstrate that certain problems – we are interested specifically in geometric vision problems – require the solution of a polynomial of a given minimum degree. As an example, we will show that the relative orientation problem for two views requires solution of a 10-th degree polynomial.

For this example, we show that one can not solve this problem by solving polynomials of degree 9 or less. We also allow the ordinary arithmetic operations, as well as extraction of radicals (square roots, cube roots, etc) up to any degree, and even Singular Value Decomposition (SVD). Still the problem can not be solved.

Our basic tool in proving our results is Galois Theory, which was invented in order to examine the question of what polynomials equations can be solved by extraction of radicals. The main study of Galois Theory is the so-called Galois group of a polynomial. If the Galois group is not solvable ([2]) then the polynomial is not solvable in terms of radicals. In particular, this holds if the

Galois group is the *symmetric group* S_n or *alternating group* A_n (defined later) for $n \geq 5$. A simple extension of this result, stated in Theorem 2.4, shows that if the Galois group of a polynomial is S_n or A_n with $n \geq 5$, then it can not be solved by extraction of radicals, or by finding the roots of a polynomial of any lower degree.

This is the theoretical basis of our non-solvability results. We explain how this may be applied to geometric vision problems. Many such problems have been solved by methods that involve the solution of a polynomial. The class of problems that can be solved in this way is quite broad, and in theory extends to the general structure and motion problem. In fact any problem whose solution involves minimizing a cost-function that is a rational expression in the problem parameters can be solved this way, since the partial derivatives of the cost function are also rational. The required solution is a point at which the partial derivatives with respect to all the variables vanish, and all such points may be found by solving a system of polynomial equations. This may be reduced to the solution of a single equation by methods such as Gröbner bases. (Of course this approach is practical only for small problems.)

To be concrete, think of the relative orientation and triangulation problems. Just because solutions exist involving polynomials of a given degree, ($n = 10$ or $n = 6$ in the problems above), does not mean that the problems can not be solved perhaps in several steps by solving polynomials of smaller degree. To show that this is in fact not possible, we examine the polynomial that arises in the problem solution, and show that it has Galois group S_n. This implies that the roots of the polynomial can not be found other than by explicitly solving a polynomial of degree n, or higher. This is not enough, however, since perhaps there is a quite different solution that involves different polynomials, or perhaps even linear techniques. Our goal is not to show that the polynomial involved in a specific solution is of a given complexity, but rather to show that the problem itself has such a complexity.

The gap is filled by showing that the roots of the polynomial involved in a specific solution are closely linked arithmetically to the numbers that appear in the solution of the problem itself. Often this is very easily shown. More precisely, we argue that the problem of finding one of the roots of a specific polynomial can be reduced to solving a given instance of the problem in question. If we can solve the problem, then we can find one of the roots of the polynomial. But, since finding one of the roots of the polynomial involves solving a polynomial of degree n, so must the problem itself. This method of reduction is illustrated in Fig 1. The terminology used in Fig 1 is explained later.

1.2 Number of Solutions and Symmetries

A measure of the degree of difficulty of a problem is the number of possible solutions it allows. However, this is not an infallible guide. Some polynomials of high degree may be solved more easily than their degree (and number of solutions) indicates. As a simple example, a polynomial $az^6 + bz^4 + d$ has degree 6, and generally 6 distinct solutions. However, we may find its roots by

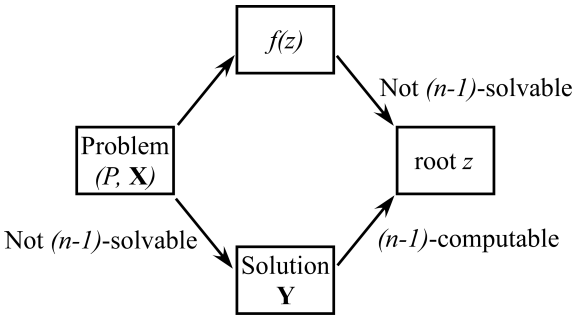

Figure 1: *An algorithm to solve a problem P with input $\mathbf{X}$ may involve the solution of a polynomial f. If the Galois group of f is S_n or A_n, with $n \geq 5$, then finding any one of the roots z of f intrinsically requires solving an n-th degree polynomial. We say f is not $\mathcal{P}_{n-1}$-solvable. Now, consider the set of numbers $\mathbf{Y}$ occurring in the solution of problem $(P, \mathbf{X})$. The reduction step is to demonstrate that z can be computed from $\mathbf{Y}$ without involving a degree n polynomial; thus z is $\mathcal{P}_{n-1}$-computable from $\mathbf{Y}$. It follows that the solution $\mathbf{Y}$ can not be $\mathcal{P}_{n-1}$ solvable given the problem instance $(P, \mathbf{X})$, otherwise z would be $\mathcal{P}_{n-1}$-solvable.*

first solving $ay^3 + by^2 + d$, and then taking square roots to find the roots z. This method avoids directly solving a 6-th degree equation. More general examples are discussed in section 4.1.

In structure from motion problems, such behaviour arises from geometric structure or symmetries specific to the problem in question. As an example of this, in the relative orientation problem, because of twisted pairs of solutions ([6]) there are actually 20 solutions for rotation ([7]). Nevertheless, solving this problem via the essential matrix requires solution of only a 10-th degree polynomial. Each essential matrix gives rise to two solutions, a twisted pair. Thus, despite having 20 solutions, this problem requires the solution of only a 10-th degree polynomial.

Another example is the three point perspective pose problem [4], which can be solved in closed form with four symmetric pairs of solutions. The symmetry corresponds to reflection of the projection center across the plane of the three points, an ambiguity that can be removed (after all the solutions have been computed) by requiring that the points reside in front of the camera. This example is particularly enlightening, because if the camera is non-central, the symmetry is no longer apparent. In this case an eighth degree polynomial can be used to solve this problem ([10]); with our method we have shown conclusively (details are omitted) that indeed an 8-th degree polynomial is required.

2 Preliminaries

All polynomials up to and including degree four are solvable in closed form (by radicals). The Greeks were able to solve the quadratic by geometric methods, see for example Euclid (325-270 BC), while formulas for the cubic and quartic were established around 1545. The quintic resisted solution and in 1824, Abel proved that the quintic is not solvable in general. Galois gave a general theory for when a polynomial is solvable in radicals.

The essence of Galois Theory is the connection between the theory of fields, particularly as it relates to solutions of polynomials, and group theory. The connection is made via the Galois group of a polynomial, or of a field extension. Essential to our approach is the ability to compute Galois groups of polynomials. To do this we use the Magma algebraic software system, [1].

We assume the reader is familiar with the basic concepts of group theory, such as *group*, *homomorphism*, *normal subgroup* and *quotient group*. In addition we assume some knowledge of field theory, including extension fields. Excellent information on these topics is available on line. We recommend the Wikipedia [14] articles on these topics which are easily found, via a web search.

Groups. We are interested in two particular groups, the symmetric group S_n, which is the group of all permutations of n symbols, and the alternating group A_n, which is the group of all *even* permutations of n symbols. Group S_n has order $n!$ and A_n has order $n!/2$. It is an important fact that for $n \geq 5$, the group A_n has no proper normal subgroups (that is, normal subgroups other than itself and the trivial group). Furthermore, S_n has only one proper normal subgroup, namely the alternating group A_n. This fact is basic to the application of Galois theory in showing the non-solvability of generic polynomial equations for degree 5 or greater. It is also the basis of our results.

Fields and field extensions. We denote the field of rationals by Q. If $x_1, \ldots, x_a$ are real or complex numbers, then $Q(x_1, \ldots, x_a)$ is the smallest field containing Q and all the x_i. We may also sometimes write $Q(\mathbf{X})$, where $\mathbf{X} = (x_1, \ldots, x_a)$. A number is in this field if and only if it may be written as $C(x_1, \ldots, x_a)/D(x_1, \ldots, x_a)$, where both C and D are multivariate polynomials over the rationals.

All fields that we consider will have characteristic zero, which simply means that they contain a copy of the integers. This assumption is harmless, and is necessary only to avoid certain technical difficulties in the next paragraph.

Given a polynomial p over a field F, we say that an extension field K of F is a *splitting field* for p if the polynomial splits into linear factors over K, but not over any smaller field. Another way of saying that K is a splitting field of some polynomial over F, is to say that K is a *finite normal extension* of F, or more briefly a *normal extension*, and denote this by $F \lhd K$. If K is an extension of a field F, we are interested in the automorphisms of K that fix every element of F. Such automorphisms form a group, known as the Galois group of the

extension K/F. If K is a splitting field of a polynomial p over F, then we also refer to this as the Galois group of the polynomial.

2.1 Definitions

Problems. We begin by defining a "problem". In our formulation, a problem is simply a mapping that takes a vector of real numbers as inputs, and produces an output vector. More general definitions are possible, but this will serve our purposes.

Definition 2.1. A *problem* is a mapping $P : \mathbb{R}^a \mapsto \mathbb{R}^b$. The problem P takes an *input vector* $\mathbf{X} = (x_1, \ldots, x_a) \in \mathbb{R}^a$ and associates to it a *solution vector* $\mathbf{Y} = (y_1, \ldots, y_b) = P(\mathbf{X}) \in \mathbb{R}^b$.

Thus, for instance in the triangulation problem, the input is a vector of numbers denoting the internal and external calibration of a set of cameras, plus a set of coordinates of corresponding image points. The solution is the vector consisting of the coordinates of the optimal $3D$ point.

In the relative orientation problem, the input consists of the coordinates of a set of matching points in two images. The solution is the vector consisting of the entries of the essential matrix (or alternatively, the entries of the rotation and translation of the relative motion). Note that this problem actually has multiple solutions. Our definition of a problem still applies; we may assume either that the mapping P arbitrarily picks one of these solutions, or provides all solutions concatenated into one vector.

A *problem instance* is a pair $(P, \mathbf{X})$, consisting of a problem and a specific input.

Classes of polynomials. We are interested in problems that can be solved by finding the roots of polynomials of a restricted kind. Most specifically, we are interested in polynomials belonging to a class, which we will denote by $\mathcal{P}_n$, consisting of

1. polynomials of degree at most n; and
2. polynomials of the form $p(z) = z^m - a$ for any m.

We note that the Galois group of a polynomial of degree at most n must be a subgroup of S_n, whereas the Galois group of the polynomial $z^m - a$ is abelian. Other wider (or more restrictive) classes $\mathcal{C}$ of polynomials are also of potential interest, as we shall see. We focus on numbers that may be computed by solving a sequence of polynomials of a given class.

Definition 2.2. Let $\mathcal{C}$ be a set of polynomials. A number y is *$\mathcal{C}$-computable* over a base field F_0, if there exists a sequence of fields $F_0 \lhd F_1 \lhd \ldots \lhd F_N$ such that $y \in F_N$ and each F_{i+1} is obtained from F_i by adjoining all the roots of some polynomial over F_i belonging to the class $\mathcal{C}$.

In this definition, we could instead have specified that each F_{i+1} is obtained by adjoining only *some* of the roots of a polynomial but it is easily seen that this is an equivalent definition. The concept of $\mathcal{C}$-computability extends also to problems, as follows.

Definition 2.3. A problem instance $(P, \mathbf{X})$ with $\mathbf{X} = (x_1, \ldots, x_a) \in \mathbb{R}^a$ is $\mathcal{C}$-solvable if each entry y_i in the solution vector $\mathbf{Y} = (y_1, \ldots, y_b)$ is $\mathcal{C}$-computable over the field $F_0 = Q(x_1, \ldots, x_a)$. A problem P is $\mathcal{C}$-solvable if every instance $(P, \mathbf{X})$ is $\mathcal{C}$-solvable for all inputs $\mathbf{X} \in \mathbb{R}^a$.

To understand this definition, note that if we start with an input vector $(x_1, \ldots, x_a)$ and apply arithmetic (addition, subtraction, multiplication or division) operations, we obtain numbers in the base field $F_0 = Q(x_1, \ldots, x_a)$. Next, by taking one or more roots of a polynomial p_0, followed by further arithmetic operations, we obtain numbers that lie in the extension field F_1. Taking roots of further polynomials, and applying further arithmetic operations extends the set of numbers that we can compute to the extension fields F_i, until eventually we reach a field in which the number y_i lies.

We will also have occasion to use terms such as $\mathcal{C}$-extension, $\mathcal{C}$-reducible and others, which involve solution of polynomials in the class $\mathcal{C}$ in a way that should be obvious from the context.

The main theorem that enables us to evaluate the degree of difficulty of a problem can now be stated.

Theorem 2.4. *Let y be a root of a polynomial p of degree $n \geq 5$ over a field F_0. If the Galois group $G(p)$ is equal to A_n or S_n, then y is not $\mathcal{P}_{n-1}$-computable over F_0.*

Although this theorem is a relatively standard application of Galois Theory, we will present a relatively complete proof so as to give the reader some feeling for why it is true.

3 Reduction

In proving that certain problems are not $\mathcal{C}$-solvable over a field F_0, our strategy is to demonstrate that some number y related to the solution of the problem is not $\mathcal{C}$-computable. This number will generally not be precisely the solution to the problem in question. However, we will be able to reduce the computation of y to solving the original problem. Thus, let $(P, \mathbf{X})$ be a problem instance and suppose that P is $\mathcal{C}$-solvable. If starting from the solution to $(P, \mathbf{X})$ we could easily compute the value y, then it would follow that y would be $\mathcal{C}$-computable. Equivalently, if we know that y is not $\mathcal{C}$-computable, then it follows that P can not be $\mathcal{C}$-solvable.

This argument can be made more formal, as follows. Given that the problem P is $\mathcal{C}$-solvable, a solution to a problem instance $(P, \mathbf{X})$ is a vector $\mathbf{Y}$ of numbers lying in an extension field F_N of $F_0 = Q(\mathbf{X})$ Now, suppose that in turn, the

number y is $\mathcal{C}$-computable over F_N, then it follows that y is $\mathcal{C}$-computable over F_0, since we can extend the field hierarchy

$$F_0 \lhd F_1 \lhd \ldots \lhd F_N$$

by a further sequence of $\mathcal{C}$-computable extensions, until ultimately we reach a field extension containing y.

We make the following definition of reducibility.

Definition 3.5. Let $\mathbf{Y} = (y_0, y_1, \ldots, y_b)$ be the solution to a problem instance $(P, \mathbf{X})$ and let $F_0 = Q(\mathbf{X})$. If a number y lies in a $\mathcal{C}$-extension of the field $F_0(\mathbf{Y}) = Q(\mathbf{X}, \mathbf{Y})$, then the problem of computing y is said to be $\mathcal{C}$-reducible to solving the problem instance $(P, \mathbf{X})$.

In other words, we can compute y starting from the inputs $\mathbf{X}$ and the solution $\mathbf{Y}$ using only arithmetic operations and solving polynomials in the class $\mathcal{C}$. Often, as in the problems considered in this paper, arithmetic operations alone suffice to compute y, and we do not need to use the input values $\mathbf{X}$.

General Strategy. The strategy for proving that a given problem P is not $\mathcal{C}$-solvable is as follows.

1. Consider a specific problem instance $(P, \mathbf{X})$ with solution $\mathbf{Y}$.
2. Find a number y with the properties that
   (a) y is not $\mathcal{C}$-computable over $F_0 = Q(\mathbf{X})$, but
   (b) y is $\mathcal{C}$-computable over $Q(\mathbf{Y})$.

It then follows that the specific problem instance $(P, \mathbf{X})$ is not $\mathcal{C}$-solvable, and hence neither is problem P. The number y mentioned here is typically a root of a polynomial arising from an algorithm used to solve the problem.

In carrying out this strategy to prove that a particular problem is not $\mathcal{C}$-solvable, it is sufficient to choose any convenient problem instance $(P, \mathbf{X})$. In practice, we choose a problem instance in which each of the components x_i of the input vector $\mathbf{X}$ is a rational number, or more usually an integer. Then, the base field $F_0 = Q(x_1, \ldots, x_a)$ is equal to the field of rationals, Q.

Later in section 8 of this paper we will consider the problem of generic inputs and show that problem instances $(P, \mathbf{X})$ are non-$\mathcal{C}$-solvable for almost all inputs $\mathbf{X}$.

4 The Theory

We require a basic result, known as the Fundamental Theorem of Galois Theory, which we will state in the following form.

Theorem 4.6 (Fundamental Theorem of Galois Theory) . *Let $F \lhd K$ be a normal field extension, and let E be an intermediate normal extension of F; thus $F \lhd E < K$. Then, there exists a homomorphism ϕ mapping $G(K/F)$ onto $G(E/F)$ with kernel $G(K/E)$. Thus*

$$\frac{G(K/F)}{G(K/E)} \approx G(E/F).$$

We will not give a complete proof of this result, but it is worthwhile to get a feeling for this theorem by outlining a proof. An element τ of $G(K/F)$ is an automorphism of K, fixing F. Its restriction to the intermediate field E induces an isomorphism of E. We need to show that τ maps E to itself, thus inducing an automorphism of E, namely an element of the Galois group $G(E/F)$. Since E is a normal extension of F it is the splitting field of some polynomial over F. Since τ fixes the base field F, it maps any root of p to some other root. Hence, τ simply permutes the roots of p, and hence maps the splitting field of p, namely E, to itself. The restriction mapping $\phi : \tau \mapsto \tau|_E$, is therefore a homomorphism of $G(K/F)$ **into** $G(E/F)$.

We consider the kernel of this mapping, in other words, what elements τ of $G(K/F)$ restrict to the identity automorphism of E. Such an element τ fixes E, and hence is an element of $G(K/E)$. Thus, $G(K/E)$ is the kernel of the homomorphism ϕ, as required.

It remains to show that ϕ maps $G(K/F)$ **onto** $G(E/F)$, namely that every automorphism of E, fixing F is the restriction of some automorphism of K. We do not prove this fact here. It depends on the assumption that K is a normal extension of F. □

We now use this theorem to prove a result about pairs of normal extensions.

Lemma 4.7. *Let F_p and F_q be normal extensions of a field F, splitting fields of the polynomials p and q respectively. Denote by F_{pq} the smallest field containing both F_p and F_q. Then, F_{pq} is a normal extension of F, and also of F_p and F_q. Moreover, $G(F_{pq}/F_p)$ is isomorphic to a normal subgroup of $G(F_q/F)$.*

The relationship between the different field extensions is as shown in the following diagram.

$$\begin{array}{ccc} F & \lhd & F_p \\ \triangle & & \triangle \\ F_q & \lhd & F_{pq} \end{array} \tag{1}$$

Proof. First, F_{pq} is a normal extension of F_p, since it is the smallest extension of F_p containing the roots of polynomial q. Thus, it is the splitting field of q over F_p. Similarly $F_q \lhd F_{pq}$. In addition, F_{pq} is the smallest field containing the roots of both p and q, hence it is the splitting field of the polynomial pq.

Now, since $F \lhd F_q \lhd F_{pq}$, according to Theorem 4.6, there is an epimorphism $\phi : G(F_{pq}/F) \to G(F_q/F)$ with kernel $G(F_{pq}/F_q)$. Also, since $F \lhd F_p \lhd F_{pq}$, according to Theorem 4.6, $G(F_{pq}/F_p)$ is a normal subgroup of $G(F_{pq}/F)$. Restricting ϕ to $G(F_{pq}/F_p)$ therefore maps $G(F_{pq}/F_p)$ onto a normal subgroup of $G(F_q/F)$.

Finally, we inquire what elements of $G(F_{pq}/F_p)$ map in this way to the identity of $G(F_q/F)$. Such an element is an automorphism of F_{pq} that fixes F_p. Since it maps to the identity in $G(F_q/F)$, it must lie in the kernel of ϕ, namely $G(F_{pq}/F_q)$. Hence, τ fixes F_q. However, since τ fixes both F_p and F_q it must fix F_{pq}, which is the smallest field containing both F_p and F_q. In other words, τ is the identity element in $G(F_{pq}/F)$. Thus the homomorphism ϕ restricted to $G(F_{pq}/F_p)$ has trivial kernel. This shows that $G(F_{pq}/F_p)$ is isomorphic to a subgroup of $G(F_q/F)$ as required. □

We now show that under certain circumstances, a field that contains one root of a polynomial must contain them all.

Theorem 4.8. *Consider a sequence of field extensions*

$$F_0 \triangleleft F_1 \triangleleft \ldots \triangleleft F_{N-1} \triangleleft F_N$$

where each F_i is a normal extension of F_{i-1}. Let p be a polynomial of degree $n \geq 5$ over F_0 with Galois group S_n or A_n. If F_N contains one of the roots of p, then it contains all the roots of p. Furthermore, if F_N is the first field in this sequence containing the roots of p, then p is irreducible over F_{N-1} and $G(F_N/F_{N-1})$ has a quotient group isomorphic to S_n or A_n.

Proof. Let $F_i(p)$ be the splitting field of the polynomial p over F_i, that is, the smallest field containing F_i and the roots of p. We have a network of field extensions of the form

$$\begin{array}{ccccccccc} F_0 & \triangleleft & \ldots & \triangleleft & F_{N-1} & \triangleleft & F_N & & \\ \triangle & & & & \triangle & & \triangle & . & \\ F_0(p) & \triangleleft & \ldots & \triangleleft & F_{N-1}(p) & \triangleleft & F_N(p) & & \end{array} \tag{2}$$

Now, starting from the left end, and applying the first part of lemma 4.7, we see that $F_i(p)/F_i$ is a normal extension, and so is $F_i(p)/F_{i-1}$ for all i. Now, according the the conclusion of lemma 4.7, we see that

$$\begin{aligned} G(F_N(p)/F_N) \quad \stackrel{\triangleleft}{\rightarrow} \quad & G(F_{N-1}(p)/F_{N-1}) \stackrel{\triangleleft}{\rightarrow} \ldots \\ & \stackrel{\triangleleft}{\rightarrow} G(F_0(p)/F_0) \stackrel{\triangleleft}{\rightarrow} S_n. \end{aligned}$$

where $A \stackrel{\triangleleft}{\rightarrow} B$ means that A is isomorphic to a normal subgroup of B. However, since the only normal subgroups of S_n are S_n, A_n or the trivial group, it follows that $G(F_N(p)/F_N)$ must be isomorphic to one of these groups. Assume now that F_N contains at least one root of polynomial p. In this case, $F_N(p)$ is actually a splitting field of a polynomial of degree at most $n-1$ over F_N, and so $G(F_N(p)/F_N)$ can not be A_n or S_n. It follows that $G(F_N(p)/F_N)$ is the trivial group, and so $F_N(p) = F_N$. Thus F_N contains all the roots of p.

Next, suppose that F_{N-1} contains no root of p. As before, $G(F_{N-1}(p)/F_{N-1})$ is isomorphic to a normal subgroup of $G(F_0(p)/F_0) \triangleleft S_n$. This time, however, $G(F_{N-1}(p)/F_{N-1})$ is not trivial, since F_{N-1} contains no

roots of p. Therefore $G(F_{N-1}(p)/F_{N-1})$ is isomorphic to A_n or S_n. It follows that p is irreducible over F_{N-1}. Finally, from the inclusion

$$F_{N-1} \lhd F_{N-1}(p) \lhd F_N$$

we deduce using Theorem 4.6 that $G(F_{N-1}(p)/F_{N-1})$ is a quotient group of $G(F_N/F_{N-1})$, as required. □

It is now possible to prove Theorem 2.4 as a corollary of Theorem 4.8.

Proof of Theorem 2.4. Let y be a root of a polynomial p of degree n over a field F_0, and let the Galois group $G(p)$ be A_n or S_n. If y is $\mathcal{P}_{n-1}$-computable over F_0, then there exists a sequence of normal extensions

$$F_0 \lhd F_1 \lhd \ldots \lhd F_N$$

where for each i, we know that $G(F_i/F_{i-1})$ is abelian, or a subgroup of S_{n-1}. However, this is incompatible with the conclusion of Theorem 4.8 that $G(F_N/F_{N-1})$ has a quotient group isomorphic to S_n or A_n.

4.1 An Example

It is instructive to give an example to show that the assumption that the Galois group be S_n or A_n is necessary in Theorem 4.8. We can not replace the condition by a condition that the polynomial p be irreducible.

Consider the polynomial $p(z) = z^4 - 2$ over the rationals, Q. This polynomial is clearly irreducible over Q. We define a sequence of splitting fields $Q \lhd Q(\sqrt{2}) \lhd Q(2^{1/4}) \lhd Q(i, 2^{1/4})$. These three extensions are splitting fields of the polynomials $z^2 - 2$, $z^2 - \sqrt{2}$ and $z^2 + 1$ respectively. Although $Q(\sqrt{2})$ does not contain any of the roots of $z^4 - 2$, the field $Q(2^{1/4})$ clearly contains the two real roots, but not the complex roots. This shows that the conclusion of Theorem 4.8 is not true for this polynomial.

Let $f(z) = z^2 - 2z - 1$ and $g(z) = z^3 - z^2 + z + 1$. It may be verified that the polynomial $p(z) = f(g(z)) = z^6 - 2z^5 + 3z^4 - 2z^3 + z^2 - 2$ is irreducible. However, this polynomial does not have a Galois group equal to S_6 or A_6, and the conclusions of Theorem 4.8 will be seen not to hold. It is possible to find the roots of the polynomial $p(z)$ in steps as follows. First, we solve $f(z)$ and get the roots $w_1 = 1 + \sqrt{2}$ and , $w_2 = 1 - \sqrt{2}$ of f. Next we solve the equations $g(z) = w_1$ and $g(z) = w_2$ to get the full set of solutions to $p(z) = f(g(z)) = 0$. In this way, we have found the roots of the polynomial $p(z)$ by solving only quadratic and cubic equations. Thus the roots of $p(z)$ are $\mathcal{P}_3$-computable.

This computation corresponds to a sequence of extensions $Q \lhd F_1 \lhd F_2 \lhd F_3$ where

1. $F_1 = Q(\sqrt{2})$ is the splitting field of f,
2. F_2 is the splitting field of $g(z) - w_1 = z^3 - z^2 + z - \sqrt{2}$ over F_1.

3. F_3 is the splitting field of $g(z) - w_2 = z^3 - z^2 + z + \sqrt{2}$ over F_2

Note that F_2 contains some but not all of the roots of $f(g(z))$. Thus, the conclusions of Theorem 4.8 are not true for this polynomial and sequence of field extensions. Neither can we conclude using Theorem 2.4 that the roots of $f(g(z))$ are not $\mathcal{P}_5$-computable. In fact, they are $\mathcal{P}_3$-computable.

The conclusions of Theorem 2.4 Theorem 4.8 do not hold here because the main hypothesis, that $p(z)$ have Galois group S_6 or A_6 does not hold. To see this, we apply the Fundamental Theorem, Theorem 4.6. Since F_1 and F_3 are both normal extensions of Q, we have

$$\frac{G(F_3/Q)}{G(F_3/F_1)} \approx G(F_1/Q) \approx Z_2 \ .$$

Here Z_2 is the group with two elements. Hence $G(F_3/Q)$ can not be A_6 (which has no non-trivial normal subgroups), and if $G(F_3/Q) \approx S_6$, then the only possibility is that $G(F_3/F_1) \approx A_6$.

However, F_3 is the splitting field of the polynomial $(g(z) - w_1)(g(z) - w_2)$ over F_1. This is a reducible polynomial over F_1 and can not therefore have Galois group A_6.

5 The Relative Orientation Problem

In previous sections we have given a method for proving that a problem can not be $\mathcal{C}$-solvable for some class $\mathcal{C}$ of polynomials. In particular, the reduction strategy described in section 3 along with Theorem 2.4 gives a general method for showing that a problem is not $\mathcal{P}_{n-1}$ computable. We now apply this method to some specific problems, starting with the relative orientation problem. The method of computation of the essential matrix followed here is based on [8].

Let $\mathbf{x}_i \leftrightarrow \mathbf{x}'_i$ be five pairs of corresponding image points. The two-view five-point calibrated relative orientation problem is to find one (or all) of the non-zero 3×3 essential matrices $\mathtt{E}$ that satisfy

$$\begin{aligned} \mathbf{x}'^{\top}\mathtt{E}\mathbf{x} &= 0 \\ 2\mathtt{E}\mathtt{E}^{\top}\mathtt{E} - \mathrm{trace}(\mathtt{E}\mathtt{E}^{\top})\mathtt{E} &= 0 \\ \det(\mathtt{E}) &= 0 \ . \end{aligned} \tag{3}$$

In general, there may be more than one essential matrix satisfying these conditions. We will show that none of them is $\mathcal{P}_n$-computable for $n < 10$. To show this, we reduce this problem to one of finding the roots of a degree 10 polynomial. We consider a specific example, defined by a set of correspondences $\mathbf{x}' \leftrightarrow \mathbf{x}$ given by

$$\begin{bmatrix} 1 & 0 & 0 \\ -1 & 1 & 0 \\ -1 & 1 & 1 \\ 1 & 1 & -1 \\ 0 & 1 & 1 \end{bmatrix} \leftrightarrow \begin{bmatrix} 1 & 0 & 0 \\ 1 & 2 & 0 \\ 1 & 0 & 1 \\ 1 & -1 & 1 \\ 1 & 1 & 0 \end{bmatrix} \tag{4}$$

$$\begin{bmatrix}
0 & 0 & -6 & 6 & -2 & -4z & 2+7z & 3z-2z^2 & -2z+4z^2 & z^2-z^3 \\
0 & 0 & -24 & -12 & -8-4z & 12-8z & -4-8z & -4z-6z^2 & 4-6z+4z^2 & -2z \\
0 & 0 & 12 & 6 & 4-6z & -12+14z & 8+10z & 2+6z-2z^2 & -4+8z+8z^2 & 2z^2 \\
0 & 6 & -6 & 24 & -2+4z & 20+18z & 8+8z & 4z+6z^2 & 12+2z+2z^2 & 2+2z^3 \\
-16 & 0 & -12 & 8 & -14z & 4-8z & 20+12z & 4-4z-6z^2 & 2z+4z^2 & -2z^2 \\
10 & 0 & -8 & 8 & 4+20z & -4-2z & -12+10z & -2+6z+8z^2 & -8z+12z^2 & -2z+2z^2+2z^3 \\
16 & 2 & 24 & 12 & 2+20z & 6+8z & 32z & 4+8z+6z^2 & 2+2z+12z^2 & 2z+2z^2 \\
-4 & -4 & 8 & -16 & -8-4z & -12-10z & -4+4z & -4z+6z^2 & 6z+2z^2 & 2z^3 \\
-6 & 4 & -4 & 14 & 4+2z & 12+12z & -4-10z & 6z+6z^2 & 8z-4z^2 & 2z+2z^2 \\
4 & 0 & -22 & -12 & 14z & 18-16z & -8z & 4z+6z^2 & 6-6z-4z^2 & 2
\end{bmatrix}
\begin{pmatrix}
x^3 \\ y^3 \\ x^2y \\ y^2x \\ x^2 \\ y^2 \\ xy \\ x \\ y \\ 1
\end{pmatrix}
= \mathbf{o}$$

Figure 2: *Matrix of equations for 5-point relative reconstruction problem.*

where the rows of the two matrices represent point correspondences, in homogeneous coordinates.

Each correspondence $\mathbf{x}'_i \leftrightarrow \mathbf{x}_i$ leads to a single linear equation in the 9 entries of $\mathtt{E}$. In all, we have 5 homogeneous linear equations in 9 unknowns. A set of four vectors can be found to span the null-space of the equation matrix, and they can be reassembled into four 3×3 matrices $\mathtt{X}$, $\mathtt{Y}$, $\mathtt{Z}$ and $\mathtt{W}$. It can be explicitly verified that a possible choice of $\mathtt{X}, \mathtt{Y}, \mathtt{Z}, \mathtt{W}$ is

$$\underbrace{\begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}}_{\mathtt{W}} \quad \underbrace{\begin{bmatrix} 0 & 0 & 0 \\ -2 & 1 & 2 \\ 0 & 1 & 0 \end{bmatrix}}_{\mathtt{X}} \quad \underbrace{\begin{bmatrix} 0 & 0 & 3 \\ 0 & 0 & 1 \\ 2 & -2 & 0 \end{bmatrix}}_{\mathtt{Y}} \quad \underbrace{\begin{bmatrix} 0 & 1 & 0 \\ 0 & 1 & 1 \\ -1 & 0 & 0 \end{bmatrix}}_{\mathtt{Z}} \tag{5}$$

To do this, we simply observe that they all satisfy the essential matrix equation $\mathbf{x}'^{\top}\mathtt{E}\mathbf{x} = 0$, and that they are linearly independent.

The essential matrix must therefore be of the form

$$\mathtt{E} = w\mathtt{W} + x\mathtt{X} + y\mathtt{Y} + z\mathtt{Z} \tag{6}$$

for some scalars w, x, y and z. The four scalars are defined only up to a common factor. The possibility $w = 0$ is tested separately and it is then assumed that $w = 1$.

Next, the non-linear constraints given in (3) are applied to the matrix $\mathtt{E}$ given by (6). This results in a set of 10 cubic equations in the unknowns x, y and z. The constraint $2\mathtt{E}\mathtt{E}^{\top}\mathtt{E} - \mathrm{trace}(\mathtt{E}\mathtt{E}^{\top})\mathtt{E} = 0$ provides 9 equations, and the constraint $\det(\mathtt{E}) = 0$ gives a single cubic equation.

Now, each of these cubic equations can be considered as a combination of the 10 monomials x^3, y^3, x^2y, y^2x, x^2, y^2, xy , x, y, 1 in x and y of degree not exceeding 3, where each monomial is multiplied by some polynomial in z. The whole set of 10 constraints may be written as a matrix equation, shown in Fig 2. Each row corresponds to a single cubic equation in x, y, z. Since this set of equations must have a non-zero solution for some value of z, the determinant of the matrix must be zero. In this example, the determinant of this matrix is

2048 $p(z)$ where

$$\begin{aligned} p(z) =& 11174859\ z^{10} + 41361525\ z^9 + 16413339\ z^8 - 91333374\ z^7 \\ & - 96079221\ z^6 + 69546666\ z^5 + 116458948\ z^4 - 26685632\ z^3 \\ & - 29121184\ z^2 - 1453312\ z - 1971200 \end{aligned} \tag{7}$$

This polynomial can be demonstrated using Magma to have Galois group S_{10}. It follows from Theorem 2.4 that the value of z obtained as a root of the polynomial $p(z)$ is not $\mathcal{P}_n$-computable for any $n < 10$. Now, looking carefully at the particular entries of the matrices `W`, `X`, `Y`, and `Z`, we see that $z/w = E_{12}/E_{33}$, and since we had normalized so that $w = 1$, we see that $z = E_{12}/E_{33}$. Therefore, the ratio E_{12}/E_{33} is not $\mathcal{P}_n$-computable, and hence neither is the essential matrix `E`.

Given the relative rotation and translation $\mathtt{R}, \mathbf{t}$ of the camera, an essential matrix `E` can be computed in closed form. Likewise, `R` and $\mathbf{t}$ can be computed from `E`. Thus, it is not important whether we consider the solution to the problem to be the essential matrix, or the motion parameters. The difficulty of the problem is the same.

6 The Triangulation Problem

Let `P` and `P'` be two 3×4 camera matrices and let $\mathbf{x} = (x, y)$ and $\mathbf{x}' = (x', y')$ be the two observed image points. The two-view L_2-optimal triangulation problem is, given $\mathtt{P}, \mathtt{P}', \mathbf{x}, \mathbf{x}'$, to find the 3D point $\mathbf{X}$ that minimizes the rational cost function that is the sum $c + c'$ of the squared reprojection errors, where

$$c = \left(\frac{(\mathtt{P}\mathbf{X})_1}{(\mathtt{P}\mathbf{X})_3} - x \right)^2 + \left(\frac{(\mathtt{P}\mathbf{X})_2}{(\mathtt{P}\mathbf{X})_3} - y \right)^2 \tag{8}$$

in the first image and analogously for the second image.

By a simple image transformation in each image that does not materially change the problem, we can assume that the two image points are both at the origin of image coordinates, namely the point with homogeneous coordinates $(0, 0, 1)$. Similarly, we may assume that the two epipoles of the cameras lie on the x-axis of the image, at points with homogeneous coordinates $(1, 0, f)$ and $(1, 0, f')$.

Since our goal is to prove that the triangulation problem can not generally be solved without solving a 6-th degree polynomial, it is sufficient to prove this fact for the particular simplified triangulation problem considered here.

In this case, the fundamental matrix has the form

$$\mathtt{F} = \begin{bmatrix} ff'd & -f'c & -f'd \\ -fb & a & b \\ -fd & c & d \end{bmatrix} \tag{9}$$

and the constants f, f', a, b, c and d are easily computed by constructing the fundamental matrix from the camera matrices, then reading them from the above form for `F`.

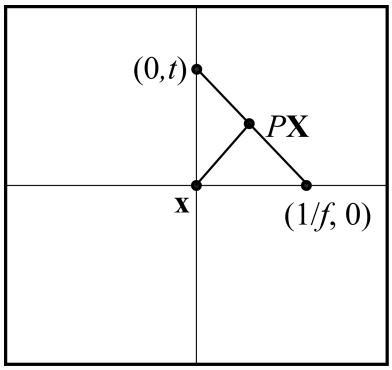

Figure 3: *The geometry of optimal 2-view triangulation. Point* $\mathbf{x}$ *is the measured image point in one image, situated at the coordinate origin. The line through the epipole* $(1/f, 0)$ *and the projection* $\mathtt{P}\mathbf{X}$ *of the optimum* $3D$ *point* $\mathbf{X}$ *meets the* y *axis at the point* $(0, t)$*, where* t *satisfies the polynomial equation (10).*

It was shown in [5] (see also [6]) that if the epipolar line in the first image corresponding to the optimal $3D$ point $\mathbf{X}$ passes through the point with homogeneous coordinates $(0, t, 1)^\top$, then t satisfies the equation

$$\begin{aligned} p(t) = t\left((at+b)^2 + f'^2(ct+d)^2\right)^2 - \\ (ad-bc)(1+f^2t^2)^2(at+b)(ct+d)\ . \end{aligned} \tag{10}$$

Note that the value t may be interpreted geometrically as the intercept of the epipolar line with the y-axis. This geometry is illustrated in Fig 3.

The polynomial in (10) is a sixth-degree polynomial. Once the roots of this polynomial are found, it is an easy matter to compute (with standard arithmetic operations) the $3D$ point that solves the triangulation. Hence, the two-view triangulation problem is generically $\mathcal{P}_6$-solvable.

The key to proving that the triangulation problem is not $\mathcal{P}_5$-solvable is to find an instance of this problem for which the polynomial p has Galois group S_6.

6.1 Non-$\mathcal{P}_5$-solvable Instance

Consider the instance of the triangulation given by the fundamental matrix (9) in which $f = f' = 1$, $a = 1$, $b = 2$, $c = 3$ and $d = 4$. Both points $\mathbf{x}$ and $\mathbf{x}'$ are at the origin. In this case, the polynomial p is $8 + 210t + 579t^2 + 612t^3 + 294t^4 + 60t^5 + 3t^6$. It may be verified that $p(t)$ is irreducible, has two complex and four real roots, and Galois group equal to S_6.

6.2 Reduction of the Problem

Finally, it is necessary to show that finding a root of this polynomial may be reduced to solving the triangulation problem. We show that from the solution

to the triangulation problem it is possible to compute the value t. Recall that we are assuming that the problem has been simplified by assuming that the measured points are at the origin, and the epipoles are on the x-axis. It is a simple matter to modify an arbitrary problem so that it is of this form.

Now, given the optimal $3D$ point $\mathbf{X}$ constituting the solution to the triangulation problem, we now project $\mathbf{X}$ into the first image, to obtain a point $\mathbf{x} = \mathtt{P}\mathbf{X}$. We also compute the epipole $\mathbf{e}$ in the first image. Next, we compute the epipolar line as the line joining $\mathbf{e}$ and $\mathbf{x}$. The intersection of this line with the y-axis is the value t.

The only operations involved in this reduction are the arithmetic field operations. Thus, computing t reduces to solving the triangulation problem.

7 Polynomials with Real Roots and SVD

Many algorithms in multiview geometry are addressed by algorithms involving the Singular Value Decomposition (SVD). Some such algorithms (for instance the Tomasi-Kanade algorithm for affine structure from motion) achieve optimal solutions using SVD. Others (such as the 8-point algorithm for two-view projective relative motion) achieve good, but non-optimal results. We are interested in the question when the SVD can lead to optimal solutions. Often algorithms involving the SVD are referred to as "linear algorithms" though this is not strictly correct, since computing an SVD can not be achieved by linear operations. We are interested in the question of whether adding the SVD to our set of available operations can make our problems solvable.

Let $\mathtt{A}$ be a matrix, and $\mathtt{A} = \mathtt{UDV}^\top$ be its SVD. The matrix $\mathtt{D}$ is diagonal, and its entries are known as the singular values of $\mathtt{A}$. Writing $\mathtt{A}^\top\mathtt{A} = \mathtt{VD}^2\mathtt{V}$, we observe that $\mathtt{VD}^2\mathtt{V}^\top$ is the eigenvalue decomposition for the matrix $\mathtt{A}^\top\mathtt{A}$, and so the singular values are the square-roots of the eigenvalues of the symmetric matrix $\mathtt{A}^\top\mathtt{A}$. Thus, finding the SVD of a matrix $\mathtt{A}$ is no harder than finding the eigenvalues of the symmetric matrix $\mathtt{A}^\top\mathtt{A}$. This is not of course a general eigenvalue problem, since the eigenvalues of a symmetric matrix are real, and in this case positive.

SVD is a weaker capability than being able to solve polynomials of arbitrary degree. In fact, it may be shown that it is weaker than being able to solve polynomials with all real roots, in that if one can solve polynomials with all real roots, then one can do Singular Value Decomposition, as expressed in the following lemma.

Lemma 7.9. *The problem of computing the Singular Value Decomposition of a matrix is $\mathcal{C}$-solvable, where $\mathcal{C}$ is the class of polynomials with all real roots.*

Proof. Given a matrix $\mathtt{A}$, form the matrix $\mathtt{A}^\top\mathtt{A}$ and extract its characteristic polynomial $p(\lambda)$. The coefficients of this polynomial are arithmetic expressions in the entries of $\mathtt{A}$. Since $p(\lambda)$ has real roots, we can find these roots, which will be real and positive. Next, taking the square root of each root (again a $\mathcal{C}$-operation) yields the singular values of $\mathtt{A}$.

The columns of the matrix $\mathtt{V}$ are the eigenvalues of $\mathtt{A}^\top\mathtt{A}$. These columns can be computed by solving the equations $(\mathtt{A}^\top\mathtt{A} - \lambda_i\mathtt{I})\mathbf{v}_i = \mathbf{0}$, where λ_i is the i-th eigenvalue of $\mathtt{A}^\top\mathtt{A}$. Since $\mathtt{A}^\top\mathtt{A} - \lambda_i\mathtt{I}$ is rank-deficient, the generator of its null space, vector $\mathbf{v}_i$ is easily computed using simple arithmetic operations. Finally, we may compute the matrix $\mathtt{U}$ from $\mathtt{U} = (\mathtt{UDV}^\top)\mathtt{VD}^{-1} = (\mathtt{A}^\top\mathtt{A})\mathtt{VD}^{-1}$. □

This theorem shows that if we can solve a problem using SVD, then we can solve it by solving polynomials with real roots. Equivalently, if it can not be solved by solving polynomials with real roots, then it can not be solved using SVD. Therefore, we concentrate on determining what problem can be solved by finding roots of polynomials all of whose roots are real. The main result in this direction is an extension of Theorem 2.4 to cover polynomials with all real roots, stated as follows.

Theorem 7.10. *Let F be a subfield of the real numbers, Suppose y is a real root of a polynomial p over F, and the Galois group of p is either A_n or S_n. Suppose in addition that the polynomial p has at least one complex root. Then y is not $\mathcal{C}$-computable over f, where $\mathcal{C}$ consists of the polynomials of degree $n-1$, polynomials of any degree with all real roots, and polynomials $z^m - a$.*

To prove Theorem 7.10 we begin by deriving a property of splitting fields of polynomials with all real roots. In the following lemma, we use the fact that if y is an element of a splitting field over a field F, then all the roots of the minimal polynomial for y over F also lie in the splitting field.

Lemma 7.11. *Let F be a sub-field of the complex numbers satisfying the condition that the real part of any element of F is also in F. Let K be the splitting field over F of a polynomial q having only real roots. Let y be a real element of K, whose minimal polynomial p over F has all real coefficients. Then all the roots of p are real.*

Proof. Since K is a splitting field of q over F, it follows that y can be written as $y = \sum f_i k_i$, where $f_i \in \mathbb{R}$ and the elements k_i are products of powers of the roots of q, and hence are real. Therefore, it follows that $y = \sum Re(f_i)k_i$, where $Re(f_i)$ is the real part of f_i. Hence, y lies in the splitting field S of q over $F \cap R$. This splitting field S must be real, since all roots of q are real. If p is irreducible over F, then it is also irreducible over $F \cap R$. Thus, all the roots of p lie in S, and hence are real. □

We are now ready to prove Theorem 7.10

Proof. If y is $\mathcal{C}$-computable, there exists a sequence of field extensions $F_0 \lhd F_1 \lhd \ldots \lhd F_N$ where each F_i is the splitting field over F_{i-1} of a polynomial in the class $\mathcal{C}$, and $y \in F_N$. Since some of the fields in this sequence may contain complex numbers, as a first step, we wish to modify this sequence of fields, replacing each field F_i in the sequence by the smallest field containing F_i and its complex conjugate field, $\bar{F}_i$. Let F_i^* be this field. We obtain a sequence of fields $F_0 \lhd F_1^* \lhd \ldots \lhd F_{N-1}^* \lhd F_N^*$. If F_i is a splitting field for polynomial q_i over

F_{i-1}, then F_i^* is a splitting field for the product $q_i\bar{q}_i$, where $\bar{q}_i$ is the polynomial whose roots are the complex conjugates of those of q_i. Thus, indeed, each of the extensions above is a normal extension. Note that if q_i is a polynomial over F_i with all real roots, then $q_i = \bar{q}_i$, so both q_i and $q_i\bar{q}_i$ have the same splitting field.

Recall that y is the root of a polynomial p, whose coefficients are in F_0 and hence are real. According to Theorem 4.8 we may assume that all the roots of the polynomial p are in F_N^*, whereas none are in the field F_{N-1}^*. Furthermore, $G(F_N^*/F_{N-1}^*)$ has quotient group isomorphic to S_n or A_n. Now, $G(F_N^*/F_{N-1}^*)$ is the splitting field of $q_N\bar{q_N}$, where q is either a polynomial of degree less than n, a polynomial $z^m - a$, or a polynomial with only real roots. In the first two cases, it is impossible for the splitting field of such a polynomial to have a quotient group isomorphic to S_n or A_n. Therefore we must conclude that F_N^*/F_{N-1}^* is the splitting field of a polynomial q_N with only real roots.

We now apply lemma 7.11 to the field extension F_N^*/F_{N-1}^* to conclude that the polynomial p has all real roots. □

The polynomials used to prove our results for the triangulation and relative motion problems also had complex roots, Theorem 7.10 applies, and we may conclude that even if SVD of arbitarily-sized matrices is allowed in addition to the other operations, the respective problems remain unsolvable, without solution of degree-6 or degree-10 polynomials respectively.

8 Genericity

Our goal has been to show that a given problem requires the solution of a polynomial of a given degree. To do this, it is sufficient to show this for a specific instance of the problem, and we choose specific examples that have the required properties. If one can not solve these specific instances without solving an n-th degree polynomial, then one can not solve the problem generally. We choose the examples for their numerical simplicity, in fact with integer data, to allow relative ease of computation of the polynomials and their Galois groups.

The reader may object that perhaps the "average" problem instance will have simpler Galois group and may be solvable by lower-degree polynomials – that in effect, the problem instances chosen are in some sense perverse. It can be shown that this is not the case. In fact, exhibiting a single example where the Galois group is S_n is sufficient to show that this is the generic case. The argument involves showing that the Galois group of the n-th degree polynomial arising from a set of data is equal to S_n, **except** on the union of a countable number of varieties in the data space, considered as a real vector space. Existence of a single example where the Galois group is S_n ensures that none of these varieties covers the whole of the input data space. Hence the set of data for which the group is not S_n has measure zero.

Consider an algorithm that takes a set of measurements and computes a result. Despite the limited number representation for a realistic computer, we will

imagine the algorithm to be computing with real numbers, and allow the inputs to the algorithm also to be real numbers. Since the inputs to the algorithm may be arbitrary, there is seemingly no restriction to the output values that may be computed by the algorithm. Given the right input, there is an algorithm that will compute any real output. The way around this slight difficulty is to consider a base field $F_0 = Q(a_0, \ldots, a_n)$ where a_1 to a_n are input values to the algorithm. Any element of the field F_0 may be computed using simple arithmetic operations on the inputs, and hence may be used as a basis for further computation. If allowable operations include more complex operations such as root-finding of polynomials, then numbers in an extension field of F_0 may be computed. For a specific algorithm we ask what extension field of F_0 the output values must lie in.

As a more specific example, we consider a polynomial $p(z) = a_0 + a_1 z + a_2 z^2 + a_n z^n$ where each $a_i \in \mathbb{R}$. We consider an algorithm that takes the values of the coefficients a_i as inputs and computes the roots of the polynomial. The outputs of the algorithm lie in the splitting field of the polynomial $p(z)$ over the base field $F_0 = Q(a_0, \ldots, a_n)$. The splitting field is $K = F(r_1, \ldots, r_n)$, where the r_i are the roots of the polynomial. The Galois group $G(K/F_0)$ may be defined in the usual way (there is no need for the base field to be the rational numbers). We now inquire when this Galois group is equal to S_n.

For polynomials of a given degree n, we may think of a vector of coefficients $\mathbf{a} = (a_0, \ldots, a_n)$ as a point in the vector space $\mathbb{R}^{n+1}$. It will be shown that vectors $\mathbf{a} \in \mathbb{R}^{n+1}$ for which the Galois group $G(K/F_0) = G(K/Q(\mathbf{a}))$ is not equal to S_n are quite scarce. In fact they lie in set of measure zero in $\mathbb{R}^{n+1}$. More exactly, there is a countable union of varieties, $\bigcup_{i=1}^{\infty} \mathcal{V}_i$ in $\mathbb{R}^{n+1}$ on which the Galois group of the polynomial fails to be S_n.

We prove our result in a sequence of lemmas. The first one involves a relationship on the roots of a polynomial over F. Central to this discussion is the concept of symmetric polynomial:

Definition 8.12. A multivariate polynomial $E(\alpha_1, \ldots, \alpha_n)$ is *symmetric* if $E(\alpha_1, \ldots, \alpha_n) = E(\alpha_{\sigma(1)}, \ldots, \alpha_{\sigma(n)})$ for every permutation σ of $1, \ldots n$.

Symmetric polynomials are related to the Galois group of a polynomial, as follows.

Lemma 8.13. *If p is a polynomial over a field F of degree n having roots $r_1, \ldots, r_n$ and Galois group not equal to S_n, then there is a* **non-symmetric** *multivariate polynomial $E(z_1, \ldots, z_n)$ over F such that $E(r_1, \ldots, r_n) = 0$.*

Proof. Note the word *non-symmetric* in the statement of the lemma. There is no difficulty in finding *symmetric* polynomials $S(z_1, \ldots, z_n)$ such that $S(r_1, \ldots, r_n) = 0$, since each ratio a_i/a_n of the coefficients of p may be expressed as a symmetric polynomial in the roots, $r_1, \ldots, r_n$. So the existence of symmetric polynomials satisfying $E(r_1, \ldots, r_n) = 0$ means nothing. It is the existence of non-symmetric polynomials that is interesting.

Let K be the splitting field of p over F. Every element of K can be expressed in the form $E(r_1, \ldots, r_n)$ for some multivariate polynomial E over F. This is

the same as saying that every element of K is a linear combination of powers of roots of p.

Arguing by contradiction, we suppose that there is no non-symmetric polynomial that vanishes on the roots of p (call this *the hypothesis*). Let σ be a permutation of $1, \ldots, n$. We define a mapping $\bar{\sigma}$ from K to K according to the following rule. If $a \in K$ and $a = E(r_1, \ldots, r_n)$, then we define $\bar{\sigma}(a) = \bar{\sigma}(E(r_1, \ldots, r_n)) = E(r_{\sigma(1)}, \ldots, r_{\sigma(n)})$.

The first thing we need to do is show that this mapping is well-defined. Namely, if

$$E_1(r_1, \ldots, r_n) = E_2(r_1, \ldots, r_n) \ ,$$

then

$$E_1(r_{\sigma(1)}, \ldots, r_{\sigma(n)}) = E_2(r_{\sigma(1)}, \ldots, r_{\sigma(n)}) \ .$$

Well, if $E_1(r_1, \ldots, r_n) = E_2(r_1, \ldots, r_n)$ then $(E_1 - E_2)(r_1, \ldots, r_n) = 0$. By the hypothesis, $E_1 - E_2$ must be a symmetric polynomial. Therefore $(E_1 - E_2)(r_1, \ldots, r_n) = (E_1 - E_2)(r_{\sigma(1)}, \ldots, r_{\sigma(n)})$, so $E_1(r_{\sigma(1)}, \ldots, r_{\sigma(n)}) = E_2(r_{\sigma(1)}, \ldots, r_{\sigma(n)})$ as required.

Next, we show that this mapping fixes points in the base field F. Thus, suppose that $a \in F$. Expressed as a polynomial expression in the roots r_i, we may choose the polynomial $E(r_1, \ldots, r_n) \equiv a$ of degree 0 over F. This choice is possible by well-definedness of the mapping. In this case, trivially $E(r_{\sigma(1)}, \ldots, r_{\sigma(n)}) = E(r_1, \ldots, r_n) = a$. Thus $\bar{\sigma}(a) = a$ as required.

Finally, we wish to show that this mapping is an automorphism of K. Let $a = E_a(r_1, \ldots, r_n)$ and $b = E_b(r_1, \ldots, r_n)$. Then $a + b = (E_a + E_b)(r_1, \ldots, r_n)$. Thus

$$\begin{aligned} \bar{\sigma}(a+b) &= \bar{\sigma}((E_a + E_b)(r_1, \ldots, r_n)) \\ &= (E_a + E_b)(r_{\sigma(1)}, \ldots, r_{\sigma(n)}) \\ &= E_a(r_{\sigma(1)}, \ldots, r_{\sigma(n)}) + E_b(r_{\sigma(1)}, \ldots, r_{\sigma(n)}) \\ &= \bar{\sigma}(a) + \bar{\sigma}(b) \ . \end{aligned}$$

Essentially the same proof shows that $\bar{\sigma}(ab) = \bar{\sigma}(a)\ \bar{\sigma}(b)$.

Let us summarize what has been shown here. Under the hypothesis that no non-symmetric polynomial vanishes on the roots of p, we deduce that there exists an automorphism of K/F that effects an arbitrary permutation on the roots of p, hence $G(K/F) = S_n$. Thus, if the Galois group is not S_n, then the hypothesis must be wrong; so the lemma is proved.

At this point, we have shown that if $p(z) = a_0 + a_1 z + \ldots + a_n z^n$ has Galois group other than S_n, then its roots satisfy a non-symmetric polynomial over $F = Q(a_0, \ldots, a_n)$. This is not quite what we want. We need to show that its coefficients (not roots) satisfy a polynomial over Q.

Lemma 8.14. *If p is a polynomial over $F = Q(a_0, \ldots, a_n)$ of degree n with Galois group not equal to S_n, then there is a non-trivial multivariate polynomial $E_Q(z_0, \ldots, z_n)$ over Q such that $E_Q(a_0, \ldots, a_n) = 0$.*

Proof. If the Galois group is not S_n, then there exists (according to lemma 8.13) a non-symmetric multivariate polynomial $E_F(z_1, \ldots, z_n)$ over F such that $E_F(r_1, \ldots, r_n) = 0$.

The base field $F = Q(a_0, \ldots, a_n)$ is not necessarily a finite (algebraic) extension of Q. However, every element of $F = Q(a_0, \ldots, a_n)$ can be written as $C(a_0, \ldots, a_n)/D(a_0, \ldots, a_n)$, where C and D are polynomials over Q. The coefficients of E_F are elements of F, so $E_F(r_1, \ldots, r_n)$ can be expressed as a rational expression in the a_i and r_i. If this is to be zero, then its numerator must be zero. This gives a vanishing polynomial in the a_i and r_i, namely

$$N(a_0, \ldots, a_n, r_1, \ldots, r_n) = 0 \ ,$$

where N is a polynomial over the rationals. Now, we add in the fact that each of a_i/a_n is an elementary symmetric expression in the roots $r_1, \ldots, r_n$. Altogether, then we have a set of $n+1$ polynomial equations in $a_0, \ldots, a_n$ and $r_1, \ldots, r_n$. We now may apply elimination theory (in particular, see theorem 2.3 on page 80 of [3]) to eliminate the r_i. This finally gives a rational polynomial expression E_Q in the coefficients a_i that will be satisfied if and only if there are a set of roots $r_1, \ldots, r_n$ satisfying $E_F(r_1, \ldots, r_n) = 0$.

Example. Consider a polynomial $a_0 + a_1 z + z^2$ where a_0 and a_1 are real numbers. As a specific simple example of the type of condition we are considering, suppose that

$$\frac{a_0 + a_1}{a_1 + 1} \ r_0 + \frac{a_0 - a_1}{a_1 - 1} \ r_1 \ = \ 0 \ . \tag{11}$$

Note that rational expressions appearing here as coefficients of the polynomial in r_1 and r_2 are elements of the field $Q(a_0, a_1)$ This can be multiplied out to give

$$(a_0 + a_1)(a_1 - 1)r_0 + (a_0 - a_1)(a_1 + 1)r_1 = 0 \ . \tag{12}$$

We add the two conditions $a_0 = r_0 r_1$ and $a_1 = -(r_0 + r_1)$, and eliminate the r_i to get a polynomial expression in a_0 and a_1. Specifically, substituting $r_1 = -(r_0 + a_1)$ in $a_0 = r_0 r_1$ and (12) gives two polynomials in r_0. Eliminating r_0 directly or by computing the resultant gives a single condition, which may be computed to be

$$4a_0^3 - 9a_0^2 a_1^2 + a_1^4 + 4a_0 a_1^4 + a_0^2 a_1^4 - a_1^6 = 0$$

This gives a polynomial over Q that must be satisfied by the coefficients a_0 and a_1 in order for (12) to be satisfied. This is the condition for the roots of a polynomial to satisfy the condition (11), and is therefore one of a countable number of conditions for the polynomial $a_0 + a_1 z + z^2$ to have Galois group not equal to S_2 over the field $Q(a_0, a_1)$.

Another way of stating the conclusion to lemma 8.14 is that the vector of coefficients $\mathbf{a} = (a_0, \ldots, a_n)$ represents a point on a variety in $\mathbb{R}^{n+1}$ defined by an $(n+1)$-variable polynomial over Q. Now, there are clearly a countable number of distinct multivariate polynomials E_Q of this kind over Q. Hence, we

deduce that the Galois group of a polynomial $a_0 + a_1 z + \ldots + a_n z^n$ is S_n except when the coefficients lie on a countable set of varieties.

We now need the following result.

Lemma 8.15. *The set of points in $\mathbb{R}^{n+1}$ that satisfy a polynomial $E(z_0, \ldots, z_n)$ is either the whole of $\mathbb{R}^{n+1}$, or else forms a set of measure zero in $\mathbb{R}^{n+1}$.*

This theorem is intuitively plausible and is a standard result about algebraic varieties. It may be proved using the Implicit Function Theorem. The proof is omitted, since it leads us too far from the main subject of this paper.

Referring to lemma 8.15, for our application, we may rule out the first case by exhibiting an example of a polynomial with Galois group S_n. Since the countable union of sets of measure zero has measure zero, we have the next step in our proof.

Lemma 8.16. *The set of points $\mathbf{a} = (a_0, a_1, \ldots, a_n)$ for which the polynomial $a_0 + a_1 z + \ldots + a_n z^n$ has Galois group not equal to S_n forms a set of measure zero in $\mathbb{R}^{n+1}$.*

Finally, we turn to the case of a particular geometric problem, such as the triangulation problem. A problem instance is defined in terms of a set of data $x_1, \ldots, x_a$ which are in general real numbers. According to our usual argument, we solve the geometric problem by forming a polynomial $p(z) = a_0 + \ldots + a_n z^n$ where the coefficients are defined in terms of the data x_i. We assume in particular that the coefficients a_i are rational expressions in terms of the x_j, thus

$$a_i = C_i(x_1, \ldots, x_a)/D_i(x_1, \ldots, x_a) \ ,$$

where both C_i and D_i are polynomials over Q. We have shown that the polynomial $p(z)$ has Galois group other than S_n only if there exists a multivariate polynomial satisfying $E_Q(a_0, \ldots, a_n) = 0$. This may be written instead as a rational expression in terms of the input data vector $\mathbf{X} = (x_0, \ldots, x_a)$ namely

$$E_Q\left(\frac{C_0(\mathbf{X})}{D_0(\mathbf{X})}, \cdots, \frac{C_n(\mathbf{X})}{D_n(\mathbf{X})}\right) \ .$$

Multiplying out the denominator finally leads to a multivariate polynomial E_D in the x_j. The polynomial $p(z)$ derived from the input data will have Galois group not equal to S_n, and hence the problem will not be $\mathcal{P}_{n-1}$-solvable, only if the data $(x_1, \ldots, x_a)$ satisfies a multivariate polynomial $E_D(x_1, \ldots, x_a) = 0$ over Q. Since there are a countable number of such polynomials, the set of input data for which the problem is $\mathcal{P}_{n-1}$-solvable forms set of measure zero in $\mathbb{R}^m$. This is true, unless one of the polynomials E_D is identically zero, which is not the case, since we have exhibited a particular example where the problem is not $\mathcal{P}_{n-1}$ solvable.

This discussion has essentially proved the following theorem.

Theorem 8.17. *Consider a problem P. Suppose there exist rational functions $a_i(\mathbf{X}) = C_i(\mathbf{X})/D_i(\mathbf{X})$ for $i = 0, \ldots, n$, where $n \geq 5$, such that for almost all input vectors $\mathbf{X}$, solving the problem instance $(P, \mathbf{X})$ is $\mathcal{P}_{n-1}$-reducible to finding the roots of the polynomial $p(z) = a_0(\mathbf{X}) + a_1(\mathbf{X})z + \ldots + a_n(\mathbf{X})z^n$. In addition, suppose that for at least one input vector $\mathbf{X}$, the polynomial $p(z)$ has Galois group S_n with $n \geq 5$. Then for almost all input vectors $\mathbf{X}$, the problem instance $(P, \mathbf{X})$ is not $\mathcal{P}_n$-solvable.*

This is the main conclusion of this section. It is likely that the theorem is also true when less restrictive relationship holds between the input data $\mathbf{X}$ and the polynomial coefficients a_i. However, as stated the theorem applies to our two example problems. We briefly demonstrate this now.

The triangulation problem. As shown in (10), the coefficients of the polynomial $p(t)$ are directly given as rational (in fact polynomial) expressions in the parameters a, b, c, d, f and f' of the fundamental matrix that describe the particular problem instance. This was assuming that the matched points were at the origin and the epipoles on the x-axis. The general case is easily tranformed to this canonical configuration so that each of these parameters is expressed simply as a rational expression in the inputs to the problem.

5-point relative orientation. This time the form of the polynomial coefficients is a little more involved. We follow the derivation of the polynomial $p(z)$ described in section 5. From 5-point correspondences one starts by defining a 5×9 equation matrix $\mathtt{A}$, the entries of which are simply products of the coefficients of the matched points (the input data). Generically, this matrix will have rank 5 and hence will have a 4-dimensional null-space, spanned by four independent vectors. These vectors may be expressed as fixed rational expressions in the entries of $\mathtt{A}$. Specifically, if we divide $\mathtt{A}$ into blocks as $\mathtt{A} = [\mathtt{B}_{5\times5}|\mathtt{C}_{5\times4}]$, then for almost all input values, $\mathtt{B}$ is non-singular and the 4 columns of the matrix

$$\mathtt{A}^{\perp} = \begin{bmatrix} \mathtt{B}^{-1}\mathtt{C} \\ -\mathtt{I}_{4\times4} \end{bmatrix}$$

are linear independent generators for the right null-space of $\mathtt{A}$. Using Cramer's rule, we can write the entries of $\mathtt{B}^{-1}$, and hence also the entries of $\mathtt{A}^{\perp}$ as rational expressions in the entries of $\mathtt{A}$.

From the columns of $\mathtt{A}^{\perp}$ we generate four matrices $\mathtt{W}$, $\mathtt{X}$, $\mathtt{Y}$ and $\mathtt{Z}$ from which the general form of the essential matrix is given by (6). As seen in section 5, the coefficients of the polynomial $p(z)$ consist of determinants of products of the entries of $\mathtt{W}$, $\mathtt{X}$, $\mathtt{Y}$ and $\mathtt{Z}$ and so are expressible as fixed rational expressions in the input data values, as required.

9 Conclusion

The method introduced in this paper effectively demonstrates that the two problems considered are optimally solved (in terms of polynomial degree) by

the existing algorithms. There is no point in searching for linear algorithms, or algorithms involving lower degree polynomials. The method is quite general and could be applied to other similar problems. As an example, we have also shown that the non-central camera pose problem requires the solution of an 8-th degree polynomial.

References

[1] The magma computational algebra system, 2006. Computational Algebra Group, University of Sydney, http://magma.maths.usyd.edu.au/.

[2] M. Artin. *Algebra*. Prentice Hall, April 1991. ISBN 0130047635.

[3] D. Cox, J. Little, and D. O'Shea. *Ideals, Varieties, and Algorithms*. Springer-Verlag, 1997. 2nd Edition, ISBN 0-387-94680-2.

[4] R. Haralick, C. Lee, K. Ottenberg, and M. Nölle. Review and analysis of solutions of the three point perspective pose estimation problem. *IJCV*, 13(3):331–356, 1994.

[5] R. I. Hartley. Kruppa's equations derived from the fundamental matrix. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(2):133–135, 1997.

[6] R. I. Hartley and N. Y. Dano. Reconstruction from six-point sequences. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages II–480 – II–486, 2000.

[7] B. K. P. Horn. Relative orientation revisited. *Journal of the Optical Society of America*, 8(10):1630–1638, 1991.

[8] H. Li and R. Hartley. Five-point motion estimation made easy. In *Proc. International Conference on Pattern Recognition*, pages 630–633, August 2006.

[9] D. Nistér. An efficient solution to the five-point relative pose problem. *PAMI*, 26(6):756–770, Jun 2004.

[10] D. Nistér. A minimal solution to the generalised 3-point pose problem. In *CVPR*, volume 1, pages 560–567, 2004.

[11] H. Stewénius, C. Engels, and D. Nistér. Recent developments on direct relative orientation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 60(4):284–294, May 2006.

[12] H. Stewénius, D. Nistér, F. Kahl, and F. Schaffalitzky. A minimal solution for relative pose with unknown focal length. In *CVPR*, San Diego, 2005.

[13] H. Stewenius, F. Schaffalitzky, and D. Nister. How hard is 3-view triangulation really. In *Proc. International Conference on Computer Vision*, pages 686 – 693, 2005.

[14] Wikipedia. www.wikipedia.org.