Using Galois Theory to Prove Structure from Motion Algorithms are Optimal

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

Richard Hartley
National ICT Australia
Australian National University
Canberra, Australia

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

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), 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 [6], and L_2-optimal two-view triangulation, which can be realized by solving a sixth degree polynomial [3]. 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 [1] 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 [8, 10, 9]. 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 ([6, 3]). We show here that to solve the 5-point relative motion problem exactly, essentially requires us 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. Because of space limitations we can only briefly 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 no higher than 9. We also allow the ordinary arithmetic operations, as well

1-4244-1180-7/07/$25.00 ©2007 IEEE

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 prove that polynomials of degree greater than 4 could not 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 has a sufficiently complex structure[1] 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, 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, and in particular the two problems that we explicitly consider in this paper have been solved by methods that involve the solution of a polynomial. The relative orientation problem can be solved using a 10-th degree polynomial, and the triangulation problem involves a 6-th degree 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. (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 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.

[1] to be precise, if the group is not solvable

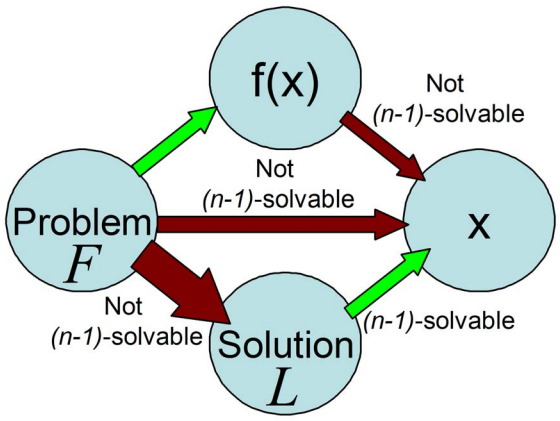

Figure 1. *An algorithm to solve a problem F may involve the solution of a polynomial f. If the Galois group of f is S_n or A_n, with $n \geq 5$, then finding the roots of f intrinsically requires solving an n-th degree polynomial (we say f is not $\mathcal{C}_{n-1}$-solvable), and neither is any root x. Now, consider the set of numbers L occurring in the solution of problem F. The reduction step is to demonstrate that x can be computed from L without involving a degree n polynomial – x is $\mathcal{C}_{n-1}$-solvable from L (green arrow). It follows that the solution L can not be $\mathcal{C}_{n-1}$ solvable given the problem instance F (red arrow), otherwise x would be $\mathcal{C}_{n-1}$-solvable.*

Genericity Our goal is 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 S_n is everywhere dense. Details of the proof are omitted for space reasons.

1.2. Numbers 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 $ax^6 + bx^4 + d$ has degree 6, and generally 6 distinct solutions. However, we may find its roots by first solving $ay^3 + by^2 + d$, and then taking square roots to find the roots x. 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 ([4]) there are actually 20 solutions for rotation and translation. 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 [2], 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 ([7]); 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, [5].

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 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. All fields that we consider will have characteristic zero, which simply means that they contain a copy of the integers[2] .

Given a polynomial p over 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". Though the following definition may not correspond to most people's conception of a problem, it focusses on the essentials for the present purposes, and gives an abstract definition of a problem as something that takes a set of inputs, and requires a solution, namely a set of numbers.

Definition 2.1. A *problem* is a mapping $P : F^a \mapsto K^b$, where F is a field (the *base field*) and K is an exten-

[2]This assumption is harmless, and is necessary only to avoid certain technical difficulties in the next paragraph.

sion field of F. The problem P takes an *input vector* $\mathbf{X} = (x_0, \ldots, x_a) \in F^a$ and associates to it a *solution vector* $\mathbf{Y} = (y_0, \ldots, y_b) = P(\mathbf{X}) \in K^b$.

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 importantly, we are interested in polynomials belonging to the following class, which we will denote by $\mathcal{C}_n$:

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

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 class of polynomials. A number y is $\mathcal{C}$*-computable* over a base field F_0, if there exists a sequence of fields $F_0 \triangleleft F_1 \triangleleft \ldots \triangleleft F_N$ such that $y \in F_N$ and each F_{i+1} is obtained from F_i by adjoining all the roots of a 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 $(P, \mathbf{X})$ is $\mathcal{C}$-solvable over a base-field F_0 if each entry y_i in the solution vector $\mathbf{Y} = (y_1, \ldots, y_n)$ is $\mathcal{C}$-computable over F_0. A problem P is $\mathcal{C}$-solvable if every instance $(P, \mathbf{X})$ is $\mathcal{C}$-solvable for all inputs $\mathbf{X} \in F_0^a$.

To understand this definition, note that if we start with a set of numbers in $F = F_0$ (consider these numbers the input data), and apply arithmetic (addition, subtraction, multiplication or division) operations, we obtain numbers in the base field F_0. 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 lies.

We will also have occasion to use terms such as $\mathcal{C}$-extension, $\mathcal{C}$-reducible and others, with meaning that should be obvious from the context.

The main theorem. 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 $\mathcal{C}_{n-1}$-computable over F_0.*

We will give a relatively complete proof of this theorem 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 be a problem and suppose that P is $\mathcal{C}$-solvable. If starting from the solution to P we could easily compute the value y, then it would follow that y would be $\mathcal{C}$-computable. Conversely, 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. A solution to a problem P over a field F_0 is a vector $\mathbf{Y}$ of numbers lying in an extension field F_N of F_0.

Now, suppose that in turn, the element 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

$$F_0 \triangleleft F_1 \triangleleft \ldots \triangleleft F_N$$

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 $\mathbf{Y} = (y_0, y_1, \ldots, y_m)$ be the solution to a problem instance $(P, \mathbf{X})$ defined over a base field F_0. If a number y lies in a $\mathcal{C}$-extension of the field $F(y_0, y_1, \ldots, y_m)$, 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 solution $\mathbf{Y}$, using only arithmetic operations and solving of polynomials in the class $\mathcal{C}$. (Often, as in the problems considered in this paper, arithmetic operations alone suffice.)

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 inputs in a base field F_0 and with solution $\mathbf{Y}$.
2. Find a number y with the properties that
   (a) y is not $\mathcal{C}$-computable over F_0.
   (b) Computing y is $\mathcal{C}$-reducible to computing $\mathbf{Y}$.

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.

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 proof of this theorem, but it is worth noting that the homomorphism ϕ mentioned in the theorem is the result of restricting an isomorphism of K/F to the intermediate field E. This provides an automorphism of E, essentially because E is a normal extension. 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*

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

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

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

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

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

where $A \stackrel{\lhd}{\to} B$ 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 $G(F_0(p)/F_0) \lhd S_n$. 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 $\mathcal{C}_{n-1}$-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.

Let $f(x) = x^2 + 2x - 1$ and $g(x) = x^3 - x^2 + x + 1$. It may be verified that the polynomial $p(x) = f(g(x)) = x^6 - 2x^5 + 3x^4 - 2x^3 + x^2 - 2$ 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 $p(x)$ in steps as follows. First, we solve $f(x)$ and get the roots $w_1 = 1 + \sqrt{2}$ and , $w_2 = 1 - \sqrt{2}$ of f. Next we solve the equations $g(x) = w_1$ and $g(x) = w_2$ to get the full set of solutions to $p(x) = f(g(x)) = 0$. In this way, we have found the roots of the polynomial $p(x)$ by solving only quadratic and cubic equations. Thus the roots of $p(x)$ are $\mathcal{C}_3$-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 $g(x) - w_1 = x^3 - x^2 + x - \sqrt{2}$ over F_1.
3. F_3 is the splitting field of $g(x) - w_2 = x^3 - x^2 + x + \sqrt{2}$ over F_2

Note that F_2 contains some but not all of the roots of $f(g(x))$. 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 $f(g(x))$ are not $\mathcal{C}_5$-computable.

5. The Relative Orientation Problem

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 $\mathcal{C}_n$-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

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

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

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

To do this, we simply observe that they each 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

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

for some scalars x, y, z and w. 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 $\mathcal{C}_n$-computable for any $n < 10$. Now, looking carefully at the particular entries of the matrices $\mathtt{X}$, $\mathtt{Y}$, $\mathtt{Z}$ and $\mathtt{W}$, 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 $\mathcal{C}_n$-computable, and hence neither is the essential matrix $\mathtt{E}$.

6. The Triangulation Problem

Let $\mathtt{P}$ and $\mathtt{P}'$ be two 3×4 camera matrices and let $\mathbf{x}$ and be $\mathbf{x}'$ 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

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

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 $\mathtt{F}$.

It was shown in [3] (see also [4]) 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.

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 $\mathcal{C}_6$-solvable.

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

6.1. Non-$\mathcal{C}_5$-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.

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

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

6.3. Singular Value Decomposition

We will summarize our results concerning Singular Value Decompostion (SVD). Many algorithms in multiview geometry are addressed by algorithms involving 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. Often algorithms involving the SVD are referred to as "linear algorithms", though this is not strictly correct. We are interested in the question of whether adding the SVD to our set of available operations can make our problems solvable.

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. It is our goal to show that we can not solve certain problems using SVD. In fact, we show the stronger result that we can not solve them *even if we can solve polynomials with all real roots*.

This conclusion relies on an extension of Theorem 2.4 to cover polynomials with all real roots, stated as follows.

Theorem 6.9. *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, where $\mathcal{C}$ consists of the polynomials of degree $n-1$, polynomials of any degree with all real roots, and polynomials $x^m - a$.*

The proof of this is a little tricky (though we have a completely written-down proof), and will not be included because of lack of space.

The polynomials used to prove our results for the triangulation and relative motion problems also had complex roots, Theorem 6.9 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.

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

[2] 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.

[3] R. Hartley and P. Sturm. Triangulation. *CVIU*, 68(2):146–157, 1997.

[4] R. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, 2000.

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

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

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

[8] 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.

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

[10] H. Stewénius, F. Schaffalitzky, and D. Nistér. How hard is 3-view triangulation really? In *ICCV*, Beijing, China, 2005.