Book Review: Algorithms in Real Algebraic Geometry by S. Basu, R. Pollack, and M.-F. Roy

Algorithms and Computation in Mathematics Series (A. M. Cohen, H. Cohen, D. Eisenbud, M. F. Singer, B. Sturmfels, eds.), vol. 10 (2$^{\underline{\text{nd}}}$ ed.), Springer-Verlag, 2006, ISBN 978-3-540-33098-1, 662 pages, $79.95

J. Maurice Rojas[*]

September 19, 2007

1 Introduction

Much like physics, the study of algorithms leads us to beautiful ideas that would otherwise be overlooked by purely abstract reflection. Moreover, as engineers and computational scientists already know, care for efficiency and effectivity in geometric constructions is not only aesthetically pleasing but ultimately useful.

Algebraic geometry, in the 20th-century sense championed by Grothendieck and his school, treats diverse ground rings on equal footing. For instance, consider the **feasibility problem over a ring** R: decide, given n-variate polynomials $f_1, \ldots, f_k$ with integer coefficients, whether there is an $x \in R^n$ such that $f_1(x) = \cdots = f_k(x) = 0$. Parallel to the blossoming of abstract algebraic geometry, the following was discovered by researchers outside, as well as inside, algebraic geometry:

1. Feasibility over $\mathbb{Z}$, in full generality, is intractable.
2. Feasibility over $\mathbb{Q}$, already for one polynomial in two variables, leads to deep open questions involving algebraic curves.
3. Feasibility over $\mathbb{R}$ and $\mathbb{C}$, and even harder problems involving quantified sentences, are theoretically tractable.

(1) is of course the celebrated solution of Hilbert's Tenth Problem, derived by Davis, Putnam, and Robinson, and completed by Matiyasevitch. (2) quickly became clear through developments of Diophantine geometry and the Birch-Swinnerton-Dyer Conjecture. (3) was first discovered by Tarski in the late 1930s, via techniques having roots in the 19$^{\underline{\text{th}}}$ century.

Beyond existential questions, it is natural to pursue geometric and topological questions: Can one somehow locate roots in a metrically precise sense? Can one compute topological invariants of real or complex zero sets? And how hard is it to answer these questions? Problems involving the approximation of real and complex solutions of systems of polynomial equations are not only mathematically interesting but of great practical importance in applications ranging from control theory to phylogenetics to semi-definite programming. Moreover, in many of these applications, one needs only information about real solutions. The need for a unified treatment of real algebraic geometry, for researchers outside as well as inside mathematics, is thus clear.

[*] Department of Mathematics, Texas A&M University TAMU 3368, College Station, Texas 77843-3368, USA, `rojas@math.tamu.edu` , `www.math.tamu.edu/~rojas` . Partially supported by NSF CAREER grant DMS-0349309, Texas MTC Grant #050245247110001, and Sandia National Laboratories.

2 The Contents

Algorithms in Real Algebraic Geometry (henceforth abbreviated ARAG) provides a self-contained treatment of some of the most important classical and modern results in semi-algebraic geometry, many authored by some subset of the trio Basu, Pollack, and Roy.

2.1 Algebraic and Topological Foundations

Chapters 1–3 start with an efficient introduction to the algebra and topology underlying real feasibility and quantifier elimination. We see the definitions of algebraic sets (the zero sets of systems of polynomial equations), constructible sets over algebraically closed fields (finite unions and intersections of algebraic sets and their complements), semi-algebraic sets (solution sets of systems of polynomial inequalities over real closed fields), and the basic notions behind quantified sentences involving polynomial inequalities. The theory of real closed fields provides a sufficiently general context to include $\mathbb{R}$, as well as Puiseux series fields (certain multivariate power series with fractional exponents) over $\mathbb{R}$. Working with general real closed fields enables the authors to make clever use of infinitesimals in their more advanced algorithms occuring later in the text. We also see, for the first time, one of the main algorithmic tools of the book: remainder sequences.

Computationally, remainder sequences (and their many variations throughout the book) are simply a structured form of univariate long division over a general coefficient ring. Thanks to work of Sturm, Sylvester, and Habicht (and more recent work by many other authors, including the authors of ARAG) it is now known that remainder sequences can be used to explicitly compute the intersection of any semi-algebraic set with any axis-parallel line segment. This central algorithmic geometric fact is first introduced in ARAG via a proof that all constructible sets are closed under projection (Theorem 1.22 of Page 25) and revisited in a proof that semi-algebraic sets are also closed under projection (Theorem 2.76 of Page 68), after the introduction of ordered, real, and real-closed fields.

At this point, ARAG has already covered an elegant proof of the Fundamental Theorem of Algebra via the theory of real closed fields, and we move on to a proof of Tarski's famous result that real closed fields admit quantifier elimination (Theorem 2.77 of Page 69). Put another way, this means that there is a finite procedure to convert any quantified sentence involving polynomial inequalities into a single formula that is true iff the original sentence was true. As an elementary example, observe that

$$\exists z_1 \in \mathbb{R}\ \exists z_2 \in \mathbb{R}\ [(az_1^2 + bz_1 + c = 0) \wedge (az_2^2 + bz_2 + c = 0) \wedge (z_2 > z_1)]$$

is true if and only if

$$(b^2 - 4ac > 0) \vee (a = b = c = 0).$$

One can of course derive this equivalence from the quadratic formula, but the power of remainder sequences lies in that they open the possibility of doing quantifier elimination in complete generality. Various closure properties of semi-algebraic sets and semi-algebraic functions are then treated in Chapter 3.

Chapter 4 sets the stage for the main effective algebraic methods of the book and, along with Chapters 8 and 10 described below, will likely be one of the most frequently consulted chapters of the book. Here we see an elegant and complete treatment of discriminants, resultants, subdiscriminants, and subresultants (each in one variable); and Sylvester-Habicht sequences and Hermite's Quadratic Form. The last result eventually yields (in Algorithm 12.7 of Chapter 12) an elegant method to count the number of solutions of a multivariate polynomial system making another polynomial positive. Chapter 4 also introduces projective space, the definition of Gröbner bases, Hilbert's Nullstellensatz (which, properly developed, allows "concise" certificates to non-solvability for polynomial systems over the complex numbers), as well as a proof of a weak version of Bézout's Theorem. It should be noted that before ARAG, a proper treatment of the topics found in Chapter 4 could only be found by combining half a dozen papers and books.

The authors have done an excellent job in finding a pedagogically satisfying order in which to introduce real algebraic geometry. By starting with an algebraic treatment of their main algorithmic tools, the authors allow the reader to see each tool twice: once in a gently abstract setting, and once again from the point of view of complexity analysis. Chapters 5 and 6 echo this pattern once again by introducing cylindrical decompositions, stratification, triangulations, as well as a short treatment of basic homology theory (including the Mayer-Viertoris Theorem and Borel-Moore homology) — each of which is revisited within the algorithms of Chapter 11 and beyond. Partioning a semi-algebraic set into managable pieces is not only useful to define homology, but also forms the foundation of important practical algorithms found later in the book.

2.2 Quantitative and Algorithmic Complexity

Chapter 7 imparts the reader with quantitative intuition on semi-algebraic sets. Indeed, common sense indicates that algorithms dealing with semi-algebraic sets should have complexity governed by their underlying topology and geometry. The theorems of Chapter 7, which include bounds on the topological complexity of semi-algebraic sets, make this sentiment rigourous. For instance, one sees that if Y is the solution set to any Boolean combination of inequalities involving at most s polynomials of degree d in k variables, then the sum of the Betti numbers of Y is bounded above by $s^2O(ds)^k$ (Theorem 7.50 has a more precise statement). In particular, this means that the preceding estimate is an upper bound on the number of real connected components of Y, and Chapter 7 thus considerably extends earlier work of Oleinik, Petrovski, Thom, and Milnor. Chapter 7 also introduces Morse theory, which is an important 20th century technique for extracting topological information about a manifold X from the critical values of a suitable real-valued function f and neighborhoods of its critical points.

Toward the middle of the book, in Chapter 8, we are introduced to algorithmic complexity through a rigourous treatment of worst-case arithmetic complexity and worst-case bit complexity. Starting with a review of some useful polynomial arithmetic and linear algebra constructions, we see explicit bounds on the number of bit operations needed to perform various low-level tasks required in algorithmic real algebraic geometry, e.g., multiplication of polynomials, the computation of characteristic polynomials, and the computation of signed subresultant sequences. It should be noted that Proposition 8.44, which gives explicit upper bounds on the bit-sizes of the coefficients of the subresultant sequence of a pair of polynomials, is centrally important in many of the complexity estimates of ARAG. In essence, Proposition 8.44 implies that coefficient growth in remainder sequences is sufficiently controllable so that one can use remainder sequences to build singly exponential algorithms for higher-dimensional problems. Furthermore, it is made clear in Chapter 8 that the bit complexity of any major algorithm from ARAG can be directly derived from the stated arithmetic complexity bounds simply by including a factor linear in the bit size of the input polynomials. (Hence the focus on arithmetic complexity through most of ARAG.) Chapter 9 then continues with explicit complexity bounds for computing quantities intimately related to counting the real roots of univariate polynomials. There we also see a treatment of the algorithmic complexity of computing the signature of a quadratic form — a task which, thanks to the development of Chapter 4, ultimately enables us to count the number of real roots of polynomial systems with finitely many complex roots.

Chapter 10 then collects various useful bounds on the size of roots of univariate polynomials, along with a novel development of the algorithmic complexity of the real roots of Bernstein polynomials. In particular, in many applications, one has polynomials expressed not as sums of monomial terms but as a linear combination of elements in a different basis. (The Bezier-Bernstein basis in geometric modelling is such a basis, and one frequently encounters the Legendre basis in the solution of differential equations.) Chapter 11 then reinforces earlier topological ideas from Chapter 5 by giving explicit algorithms for the cylindrical algebraic decomposition introduced earlier. As a consequence, we see our first complexity bound for quantifier elimination: doubly exponential in the number variables (Algorithm 11.15 and Exercise 11.8). Moreover, we see a complexity lower bound via the existence of sentences for which this doubly exponential dependence is unavoidable (Theorem 11.18). Note, however, that this need not obstruct faster algorithms in special settings. Finding such speed-ups remains a highly active area of research. Among many other geometric algorithms, we also see an algorithm for determining the topology of any (even singular) real algebraic plane curve (Algorithm 11.18).

As we have begun to see, many of the algorithms of ARAG reduce — at a low level — to computing remainder sequences. However, one also sees a higher-level problem occuring repeatedly: the solution of polynomial systems. Chapter 12 thus provides an exposition on some of the most important techniques for this central problem. In particular, the authors complete their description of Gröbner bases, and move on to the explicit construction of the multiplication map in a polynomial quotient ring. (From the latter, one can employ Hermite's quadratic form to count real roots and more.) Chapter 12 also gives a complete discussion of rational univariate reduction, expanding on the algebraic framework based on idempotents introduced earlier in Chapter 4. For instance, the following example gives a taste of Sections 12.4 and 12.6: Consider the following system of equations:

$$F = (1 + 2x - 2x^2y - 5xy + x^2 + 3x^3y, 2 + 6x - 6x^2y - 11xy + 4x^2 + 5x^3y).$$

(This example was specifically constructed to have a real zero set Z consisting of two isolated points and a vertical line: $\left\{(1,1), \left(\frac{1}{7}, \frac{7}{4}\right), \{-1\} \times \mathbb{R}\right\}$.) The techniques of Chapter 12 then allow us to compute the following three polynomials

$$h(t) = -153 + 120t + 1540t^2 + 1600t^3 + 448t^4$$
$$h_1(t) = -\frac{11762}{7511} + \frac{19150}{22533}t + \frac{114736}{22533}t^2 + \frac{7264}{3219}t^3$$
$$h_2(t) = -\frac{5881}{7511} + \frac{32108}{22533}t + \frac{57368}{22533}t^2 + \frac{3632}{3219}t^3,$$

which encode a finite set of points intersecting each and every real connected component of Z. In particular, evaluating (h_1, h_2) at the roots of h indeed recovers the isolated points $\left\{(1,1), \left(\frac{1}{7}, \frac{7}{4}\right)\right\}$, as well as the pair of points $\left\{(-1,1), \left(-1, \frac{1}{4}\right)\right\}$ which lie in the sole 1-dimensional component of Z.

Thanks to the authors' initial algebraic framework, they are able to employ infinitesimals to circumvent general position assumptions (needed in certain intermediate projection constructions) and perform the difficult tasks of Chapter 12 in complete generality. Sampling points in the connected components of real algebraic sets is one such example, and toward the end of the chapter we see another example through an algorithm for computing the Euler-Poincaré characteristic of an algebraic set over a real closed field.

Chapters 13 and 14 then revisit the quantifier elimination results introduced in Chapters 2 and 11. In particular, the exposition begins by detailing how to find all sign conditions realized by a set of real polynomials, and even sample points realizing all such sign conditions. As a consequence, we see an algorithm for deciding whether the real zero set of a single polynomial of degree d in k variables is zero-dimensional, within arithmetic complexity $d^{O(k)}$ (Theorem 13.17). We also see an improved algorithm for quantifier elimination: Algorithm 13.13, among other things, allows one to decide existential sentences involving inequalities in s polynomials of degree $\leq d$ in $\leq k$ variables within $s^k d^{O(k)}$ arithmetic operations. In particular, the complexity is singly exponential in the number of variables, and Chapter 14 gives a similar complexity bound for sentences involving a fixed number of quantifier alternations, and also gives algorithms for minimizing polynomials over arbitrary semi-algebraic sets (also finding a minimizer should one exist) and computing the dimension of arbitrary semi-algebraic sets. Chapter 13 also makes the important observation (in Remark 13.10, Page 513) that most of the singly exponential algorithms of ARAG use a polynomially bounded amount of space and can thus be placed within the complexity class **PSPACE**. In particular, this implies that most of the algorithms of ARAG (for instance, those for the existential theory of the reals) can be efficiently parallelized.

Chapters 15 and 16 then conclude ARAG with an algorithm to derive a very special (semi-algebraic) 1-dimensional subset of any input semi-algebraic set, called a **roadmap**, from which one can answer even harder topological queries. For instance, an important problem occuring in robotics is the famous **piano mover's problem**: how does one find a path through a set of (semi-algebraic) obstacles? (It is not hard to show that trying to move a semi-algebraic piano through a semi-algebraic obstacle course is reducible to the special case of connecting two points in a semi-algebraic set.) The roadmap was introduced by John Canny in the late 1980s and in Chapter 16 we see how the roadmap can be used to compute descriptions of all connected components of a semi-algebraic set (Algorithm 16.8) and compute the first Betti number of any semi-algebraic set (Algorithm 16.16).

2.3 Remarks on the Writing

As the authors clearly state in their introduction, there are recent topics that, even in a book this size, can not be addressed in a complete way. However, even for the reader interested in more advanced topics (e.g., the recent connections between semi-definite programming and real algebraic geometry [Par03] or alternative approaches to the existential theory of the reals which incorporate numerical conditioning and numerical stability [CS99]), ARAG provides important foundational complexity results, some already optimal, to which more advanced algorithms can be compared.

The only criticism of ARAG this reviewer can supply is thus minor and non-mathematical: There are a surprisingly high number of grammatical and typographical mistakes. (Without trying too hard, I recorded over 117 such mistakes throughout the book.) Indeed, this is the first time I have seen a Springer-Verlag book with such sloppy proof-reading. Nevertheless, even this many typos do not detract from the clear and excellent coverage ARAG provides for students and researchers alike.

3 Why you should buy the book under review

The authors have clearly done a tremendous service by providing a single, self-contained, and surprisingly complete source for the foundations of algorithmic real algebraic geometry. They have also organized their material in a way that can be reasonably taught to graduate students. There are simply too many natural quantitative and algorithmic questions that go completely untouched in classic, and even recent texts on algebraic geometry. To at last find a single book that finally answers these questions is a beautiful surprise.

References

[CS99] Cucker, Felipe and Smale, Steve, *"Complexity estimates depending on condition and round-off error,"* J. ACM 46 (1999), no. 1, pp. 113–184.

[Par03] Parrilo, Pablo A., *"Semidefinite programming relaxations for semialgebraic problems,"* Algebraic and geometric methods in discrete optimization, Math. Program. 96 (2003), no. 2, Ser. B, pp. 293–320.