3

Quantifier elimination and projection

The circle $\{ x^2 + y^2 = 1 \}$ is an algebraic variety and over $\mathbb{C}$ Chevalley's Theorem predicts that its projection onto the x-axis is constructible, i.e., a finite boolean combination of varieties. But over the real numbers, the projection is the interval $[-1, 1]$ which cannot be expressed as such a combination — adding the order relation to our vocabulary seems to be necessary to describe the image of the projection. This, in turn, broadens the sets we can consider projections of. Now, are even more relations needed to understand the projections of semialgebraic sets? The answer given in this chapter is: no, semialgebraic sets are closed under projection.

3.1 Quantifier elimination

The main result of this chapter is stated in model-theoretic language (see Appendix B), and its proof follows a model-theoretic approach [MT03, Section 2.5], because in this way it attains its strongest form. Briefly, it states that every first-order formula involving quantifiers is equivalent, over every real-closed field, to a quantifier-free formula (not depending on the field). The depth of this formalistic statement may not seem obvious right now, but in the subsequent sections we derive from it some of the most celebrated geometric results in real algebraic geometry.

Theorem 3.1. The theory of real-closed fields admits quantifier elimination in the language of ordered rings.

Lemma 3.2. To prove quantifier elimination, it suffices to consider formulas of the form $\exists v : f(v, w) = 0 \wedge \bigwedge_j g_j(v, w) > 0$ for $f, g_j \in \mathbb{Z}[v, w]$.

Proof. Let ϕ be any first-order formula in the language of ordered rings. By an easy result in first-order logic, we may assume that ϕ is in *prenex form*, i.e., $\phi = Q_n v_n \dots Q_1 v_1 : \phi'$ where each $Q_i \in \{ \exists, \forall \}$ and ϕ' is quantifier-free. This quantifier-free formula, where all the v_i are regarded as free variables, defines a semialgebraic set $K \subseteq \mathbb{F}^N$ over any real-closed field $\mathbb{F}$, where $N \geq n$ is the number of free variables in ϕ'. The proof of Lemma 2.5 gives a purely formal method of transforming ϕ' into an equivalent formula of the form

$$\bigvee_i \left(\bigwedge_j f_{ij} = 0 \wedge \bigwedge_k g_{ik} > 0 \right)$$

for integer polynomials $f_{ij}, g_{ik} \in \mathbb{Z}[v_1, \ldots, v_n]$. Since $\bigwedge_j f_{ij} = 0$ is equivalent to $\sum_j f_{ij}^2 = 0$, we may suppose that there is only one equation $f_i = 0$ per disjunction term.

We eliminate quantifiers inductively from the inside out. Since $\forall v : \phi'$ is equivalent in $\mathbb{F}$ to $\neg\exists v : \neg\phi'$, it suffices to eliminate existential quantifiers. Furthermore, since the formula $\exists v : \bigvee_i \phi_i'$ is equivalent to $\bigvee_i (\exists v : \phi_i')$, we arrive at the special case $\exists v : f(v, w) = 0 \wedge \bigwedge_j g_j > 0$. □

The previous lemma reduces general formulas to polynomial sign constraints. The task of eliminating $\exists v$ boils down to characterizing those regions in which a given polynomial has a fixed sign.

Lemma 3.3. Fix a degree d. We identify a polynomial $f = \sum_{i=0}^{d} f_i x^i \in \mathbb{F}[x]$ of degree at most d with its coefficient vector $(f_0, \ldots, f_d)$. There exist quantifier-free formulas $\mathsf{nroots}^d(f)$ and $\mathsf{root}_k^d(f, \rho)$ such that over every real-closed field $\mathbb{F}$:

1. $\mathsf{nroots}_p^d(f)$ is true in $\mathbb{F}$ if and only if f has exactly p roots in $\mathbb{F}$.
2. $\mathsf{root}_k^d(f, \rho)$ is true in $\mathbb{F}$ if and only if ρ is the k^{th} root of f in $\mathbb{F}$.
3. nroots_p^d and root_k^d do not depend on $\mathbb{F}$.

(For our purposes here, the zero polynomial counts as not having any root.) It follows that for every fixed d and p there is a quantifier-free formula characterizing the non-zero vectors $(f_0, \ldots, f_d)$ whose associated polynomial has exactly p roots.

This lemma already hints at some of the surprising results we will get out of this investigation: since nroots_p^d is independent of the field $\mathbb{F}$, a univariate polynomial has the same number of roots no matter which real-closed field it is considered over, proving Lemma 1.19 we used earlier without proof.

Proof. We give algorithms that solve the *root counting* and *root isolation* problems in real algebraic geometry. These algorithms are primitive recursive and only make use of the language of ordered rings. This makes it clear how to write out a quantifier-free formula describing the result of the computation. We provide these algorithms by induction (primitive recursion) on d. For $d = 0$, there is no root since $f_0 \neq 0$. Suppose the formulas for all degrees $< d$ have been constructed and consider degree d. We may assume that $f_d \neq 0$ (because the condition $f_d = 0$ can be detected in a quantifier-free way and then the problem is solved by the induction hypothesis). We can write down the derivative $f' = \sum_{i=1}^{d} i f_i x^{i-1}$ which is non-zero. Hence we can count and list roots of f' by the induction hypothesis.

Let $\rho_1 < \cdots < \rho_p$ be the distinct roots of f'. By Corollary 1.23, f' has constant sign in every interval $I_i = (\rho_i, \rho_{i+1})$ which shows, by Exercise 3.4, that f is monotone on these intervals. Now consider $f(\rho_i)$ and $f(\rho_{i+1})$:

(1) If both values are zero, then Rolle's theorem (Exercise 3.3) shows that f' has a root in I_i which contradicts the completeness of our listing of roots.

(2) If they are non-zero and have the same sign, then by monotonicity there is no root of f in I_i. Similarly, if exactly one of the endpoints happens to be a root of f, then there is no root in I_i but one on the boundary.

(3) If they have opposite signs, there must be a root of f in I_i by Corollary 1.23. There cannot be two roots by Rolle's theorem.

This analysis shows how to detect whether I_i contains a root or not. If it does, the root is unique and I_i is an isolating interval. This allows us to define this root uniquely by the formula $f(x) = 0 \wedge \rho_i < x < \rho_{i+1}$.

We have the intervals $I_0 = (-\infty, \rho_1)$ and $I_p = (\rho_p, \infty)$ left to consider. Again, f' has constant sign on them and f is monotone so it has at most one root. Lagrange's bound (Exercise 3.5) shows that a polynomial f has no root outside of the closed interval $[-R, R]$ where $R = \sum_{i=0}^{d} |f_i/f_d|$ which reduces these infinite intervals to the bounded setting considered before. (Note that the bound $R(f) \geq R(f')$, so $[-R, R]$ also encloses all roots of f'.)

This gives an algorithm to count the number of roots of f and to compute isolating intervals for them based on computing roots for f' and checking signs of f on these roots. For every fixed degree d, this algorithm can be formalized as a quantifier-free formula in the language of ordered rings. □

Lemma 3.4. To prove quantifier elimination, it suffices to consider formulas of the form $\exists v : \bigwedge_j g_j(v, w) > 0$ for $g_j \in \mathbb{Z}[v, w]$.

Proof. By Lemma 3.2 it is enough to eliminate the quantifier from the formula $\exists v : f(v, w) = 0 \wedge \bigwedge_j g_j(v, w) > 0$ for $f, g_j \in \mathbb{Z}[v, w]$. Write $f(v, w) = \sum_{k=0}^{d} f_k(w) v^k$ for $f_k \in \mathbb{Z}[w]$. Now the existence of a solution v to the polynomial system $\{ f = 0, g_j > 0 \}$ is equivalently expressed as

$$\left(\bigwedge_{k=0}^{d} f_k(w) = 0 \wedge \exists v : \bigwedge_j g_j(v, w) > 0 \right) \vee$$

$$\left(\bigvee_{1 \leq k \leq p \leq d} \mathsf{nroots}_p^d(f(-, w)) \wedge \mathsf{root}_k^d(f(-, w), \rho_{pk}) \wedge \bigwedge_j g_j(\rho_{pk}, w) > 0 \right),$$

in which the latter disjunctive term is quantifier-free by Lemma 3.3. The quantifier has moved inwards to a system of strict polynomial inequalities, as claimed. □

Finally we are ready to prove the main result:

Theorem 3.1. By Lemma 3.3 we can describe for every $g_j(v, w)$ the vectors b such that $g_j(-, b)$ has exactly p_j roots and to list these roots $\rho_1^j(b) < \cdots < \rho_{p_j}^j(b)$ as functions of b, using quantifier-free formulas. Again, this defines a partition of the underlying field into intervals $I_i^j(b) = (\rho_i^j(b), \rho_{i+1}^j(b))$ and $I_0^j(b) = (-\infty, \rho_1^j(b))$ and $I_{p_j}^j(b) = (\rho_{p_j}^j(b), \infty)$ on which g_j has constant sign. This sign can be determined in a uniform fashion by evaluating g_j on a point in the interior of these intervals.

There are finitely many possible orderings of the roots ρ_i^j of all the g_j; each such ordering shatters the underlying field into finitely many intervals on which the signs of all the g_j are fixed. A v such that $g_j(v, w) > 0$ exists if and only if there is one interval among them on which the sign of all g_j is positive. □

Corollary 3.5. The theory RCF of real-closed fields is complete and decidable.

Proof. By Theorem 3.1 every first-order sentence is equivalent to a quantifier-free sentence. But a sentence without quantifiers in the language of ordered rings is a boolean combination of equations and inequalities in explicit integers. Either this sentence or its negation is true in every ordered field which proves completeness. The proof of Theorem 3.1 describes a symbolic algorithm to compute the quantifier-free sentence and this proves decidability. □

Corollary 3.6. A set $K \subseteq \mathbb{F}^n$ over a real-closed field $\mathbb{F}$ is definable in the language of ordered rings if and only if it is semialgebraic.

Proof. Clearly every semialgebraic set is definable. Let $K = \{ x \in \mathbb{F}^n : \mathbb{F} \models \phi(x, b) \}$ be an arbitrary definable set with a first-order formula ϕ and $b \in \mathbb{F}^m$. By Theorem 3.1 there is an equivalent quantifier-free formula ψ defining this set. But without quantifiers, ψ consists of a boolean combination of polynomial equations and inequalities with integer coefficients and parameters $b \in \mathbb{F}^m$. This is a semialgebraic description of K. □

Remark 3.7: Algorithmic aspects. Quantifier elimination in practice does not follow the algorithm outlined in the proof of Theorem 3.1 but instead relies on *cylindrical algebraic decomposition (CAD)*. With this method, a semialgebraic set is first decomposed into special cells which simplify the elimination of existential quantifiers. Still, state of the art quantifier elimination algorithms have a running time which is doubly exponential in the number of quantifiers [BPR06, Section 11.3].

However, if there are no quantifier alternations, say when solving the *existential theory of the reals*, i.e., the existence of points in semialgebraic sets defined by quantifier-free formulas, there exist algorithms which eliminate the quantifier block in single exponential time in the number of quantified variables and using only polynomial space; see [BPR06, Chapter 13].

This shows that the theory of real-closed fields consists of all true sentences in the language of ordered rings about the field $\mathbb{R}$ with its usual order. Since all ordered fields include $\mathbb{Q}$, all real-closed fields include $\mathsf{rcl}(\mathbb{Q}) = \mathbb{R} \cap \mathsf{acl}(\mathbb{Q})$, the field of *real algebraic numbers*. Unlike $\mathbb{R}$, this field is amenable to symbolic computations since every finite collection $\alpha_1, \ldots, \alpha_k \in \mathsf{rcl}(\mathbb{Q})$ is algebraic over $\mathbb{Q}$ and hence by Theorem A.1 they are all contained in a finite primitive extension $\mathbb{Q}(\alpha) = \mathbb{Q}[x]/\langle f \rangle$ for a monic polynomial f with integer coefficients. In this representation, the field arithmetic can be implemented exactly on a computer using arbitrary-length rational numbers.

Example 3.8. Quantifier elimination for real-closed fields is implemented Wolfram Mathematica [WM] and the free software QEPCAD [Bro03]. To see a semialgebraic description of all coefficient vectors (a, b, c) for which the quadratic equation $ax^2 + bx + c$ has a real solution:

```
Simplify @ Reduce[Exists[x, a x^2 + b x + c == 0], Reals]
```

$$(b = 0 \wedge ((c > 0 \wedge a < 0) \vee (a > 0 \wedge c < 0))) \vee (b \neq 0 \wedge 4ac \leq b^2) \vee c = 0$$

As we know, this hinges on the discriminant $b^2 - 4ac$ being non-negative. Quantifier elimination arrives at the same result automatically and points out all the edge cases as well. △

3.2 Projection of semialgebraic sets

Tarski–Seidenberg theorem. Let $K \subseteq \mathbb{R}^{n+m}$ be semialgebraic and $\pi : \mathbb{R}^{n+m} \to \mathbb{R}^m$ the projection deleting the last m coordinates. Then the image $\pi(K)$ is semialgebraic.

Proof. Let $K = \{ (x, y) \in \mathbb{R}^{n+m} : \mathbb{R} \models \phi(x, y, b) \}$ be a semialgebraic set written as a definable set with parameters $b \in \mathbb{R}^p$. The projection is described by

$$\pi(K) = \{ x \in \mathbb{R}^n : \mathbb{R} \models \exists y : \phi(x, y, b) \}.$$

This set is definable and hence semialgebraic by Corollary 3.6. □

Corollary 3.9. Let $f : X \to Y$ be a semialgebraic map between semialgebraic sets X and Y. Then its image $\operatorname{im}(f) \subseteq Y$ is semialgebraic.

Proof. Since f is semialgebraic, its graph $\Gamma(f)$ is a semialgebraic set in $\mathbb{R}^n \times \mathbb{R}^m$. The image of f is the projection of $\Gamma(f)$ deleting the first n coordinates, which is semialgebraic by the Tarski–Seidenberg theorem. □

Corollary 3.10. Let $K \subseteq \mathbb{R}^n$ be semialgebraic. Then its topological closure $\overline{K}$, its interior K° and its boundary ∂K are semialgebraic.

Proof. Let $K = \{ x \in \mathbb{R}^n : \mathbb{R} \models \phi(x, b) \}$. By definition the closure is

$$\overline{K} = \left\{ x' \in \mathbb{R}^n : \mathbb{R} \models \forall \varepsilon \, \exists x : \varepsilon > 0 \Rightarrow (\phi(x, b) \wedge \|x' - x\|^2 < \varepsilon) \right\}.$$

Eliminating the two quantifiers gives a semialgebraic description of $\overline{K}$. Similarly, the interior is

$$K^\circ = \left\{ x' \in \mathbb{R}^n : \mathbb{R} \models \exists \varepsilon \, \forall x : \varepsilon > 0 \wedge (\|x' - x\|^2 < \varepsilon \Rightarrow \phi(x, b)) \right\}$$

and the same argument applies. The boundary is semialgebraic as the difference of the two semialgebraic sets. □

If K is defined by a formula ϕ with parameters b, we will use the expression $\exists x \in K : \psi$ as a short-hand for $\exists x : \phi(x, b) \wedge \psi$. Abbreviations like $\exists \varepsilon > 0 : \psi$ should be understood in the same way.

A set is closed if and only if it equals its closure, and open if and only if it equals its interior. By Corollary 3.10 and Corollary 3.5 this gives an algorithm to decide if a $\mathbb{Z}$-defined semialgebraic set is open, closed or neither.

Proposition 3.11. Let $K \subseteq \mathbb{R}^n$ be a locally closed semialgebraic set. Then K is semialgebraically homeomorphic to a closed set in $\mathbb{R}^{n+1}$.

Proof. As it is locally closed, K is the intersection of a closed set C with an open set U in $\mathbb{R}^n$. We can take $C = \overline{K}$ and $U = (\mathbb{R}^n \setminus C) \cup K = \mathbb{R}^n \setminus (\overline{K} \setminus K)$, so both C and U can be supposed to be semialgebraic. If A is already closed, then it can be embedded as $A \times \{ 0 \}$ and we are done.

Otherwise $K' = \mathbb{R}^n \setminus U$ is non-empty and we consider the function $\mathbb{R}^n \ni x \mapsto \operatorname{dist}(x, K') = \inf \{ \|x' - x\| : x' \in K' \}$. This function is well-defined since K' is non-empty. By the triangle inequality for $\|-\|$, it is continuous. That its graph is semialgebraic is seen in a similar manner to Corollary 3.10. Clearly $\operatorname{dist}(x, K')$ is non-negative and it vanishes if and only if $x \in \overline{K'} = K'$. Note in particular that $\operatorname{dist}(x, K')$ is positive on K. We use this function to map

$$K \ni x \mapsto (x, 1/\operatorname{dist}(x, K')).$$

This map is clearly injective, semialgebraic and continuous. Therefore it is a semialgebraic homeomorphism onto its image which is the closed semialgebraic set

$$\left\{ (x, d) \in \mathbb{R}^{n+1} : x \in \overline{K} \wedge d \operatorname{dist}(x, K') = 1 \right\}. \qquad \square$$

3.3 Exercises

Choose exercises to solve from the list below. The target value this week is 15 points. By solving more exercises, you can get up to 20 points. Solutions must be submitted on MyCourses by **Friday, May 19, 12:00**. Use of computer algebra software like `Mathematica` is highly encouraged. Also submit the source code for your computations.

3.1 Give an axiomatization of real-closed fields in the language of ordered rings. 3 points

3.2 Give a quantifier-free formula (i.e., an isolating interval) for the second root in $\mathbb{R}$ of the polynomial $x^2 - \sqrt[3]{2}x - 1/2$. 3 points

3.3 Prove Rolle's theorem: if $f \in \mathbb{F}[x]$ over a real-closed field $\mathbb{F}$ has $f(a) = f(b)$ with $a < b$, then f' has a root strictly between a and b. 2 points

3.4 Let $f \in \mathbb{F}[x]$, $\mathbb{F}$ real-closed, and suppose that f' has constant sign on the interval (a, b). Prove that f is monotone on (a, b). 3 points

3.5 Prove Lagrange's bound: a polynomial $f = \sum_{i=0}^{d} f_i x^i \in \mathbb{F}[x]$ of degree d has no root outside $[-R, R]$ where $R = \sum_{i=0}^{d} |f_i/f_d|$. 3 points

3.6 Show that semialgebraic isomorphism is an equivalence relation. 2 points

3.7 Prove that the set of coefficients (a, b, c, d, e) for which the quartic equation $ax^4 + bx^3 + cx^2 + dx + e$ has four distinct real solutions is semialgebraic. Can you give a semialgebraic description? 4 points

3.8 Two vectors $(a, b), (c, d) \in \mathbb{R}^2$ are linearly dependent if there exist $(\lambda, \mu) \neq (0, 0)$ such that $\lambda(a, b) + \mu(c, d) = (0, 0)$. Find an equivalent condition for linear dependence by eliminating quantifiers in the previous sentence. 2 points

3.9 Consider an ellipse $E = \{ ax^2 + by^2 = c \}$ with $a, b, c > 0$ and a hyperbola $H = \{ xy = d \}$ in $\mathbb{R}^2$. For generic complex choices of these parameters, there are exactly four complex intersection points. Find a semialgebraic characterization of the parameters $(a, b, c, d) \in \mathbb{R}^4$ for which all four intersection points are real. 4 points

3.10 Consider the cubic curve $C = \{ x^3 + x^2y - xy^2 = 10 \} \subseteq \mathbb{R}^2$. Find a semialgebraic description of all ellipses $E = \{ ax^2 + by^2 = c \}$ with $a, b, c > 0$ which touch C in at least one real point. Draw the cubic and such an ellipse. Can you find an ellipse with three distinct contact points? (E *touches* C in a so-called *contact point* (x, y) if this point lies on $E \cap C$ and the gradients of E and C at (x, y) are linearly dependent.) 6 points

3.11 Let K be a $\mathbb{Z}$-defined semialgebraic set in $\mathbb{R}^n$. Prove that every isolated point of K has algebraic coordinates over $\mathbb{Q}$. (An *isolated point* of a semialgebraic set K is one which has an open neighborhood not containing any other point of K.) 4 points

3.12 The *Hausdorff distance* between two semialgebraic sets $X, Y \subseteq \mathbb{R}^n$ is $h(X, Y) := \inf \{ \varepsilon \geq 0 : X \subseteq Y_\varepsilon \wedge Y \subseteq X_\varepsilon \}$ where $X_\varepsilon = \bigcup_{x \in X} B_\varepsilon(x)$ is the ε-thickening of X. Show that: (1) If Y is fixed, then $h(x, Y) := h(\{ x \}, Y)$ is a semialgebraic function of $x \in \mathbb{R}^n$. (2) If X and Y are $\mathbb{Z}$-defined semialgebraic sets, then the Hausdorff distance is computable and the result is an algebraic number. 8 points