Part 1: Homotopies and the Fundamental Theorem of Algebra

Tim Duff

February 3, 2025

Let $\mathbb{K}$ be a field. Recall that $\mathbb{K}$ is said to be *algebraically closed* if every nonconstant univariate polynomial with coefficients in $\mathbb{K}$ has a root. For example, neither the field of real numbers $\mathbb{R}$ nor its subfield consisting of rational numbers $\mathbb{Q}$ is algebraically closed, since the polynomial x^2+1 has no real roots.

Here is one reason why mathematicians love the field of complex numbers $\mathbb{C}$.

Theorem 1 (Fundamental Theorem of Algebra)**.** $\mathbb{C}$ is algebraically closed.

It's likely you've seen this theorem stated in a high school algebra course. I will explain a proof of Theorem 1 that will hopefully help you understand *why* it is true. The proof is essentially constructive, and introduces the main ideas behind *homotopy continuation*, which can be used to numerically solve systems of equations in more than one variable. The key idea is to think of a univariate polynomial $a_d x^d + a_{d-1}x^{d-1} + \cdots + a_0 \in \mathbb{C}[x]$ as a point $(a_d, \ldots, a_0) \in \mathbb{C}^{d+1}$ in a *space of problems.* Within this space, most problems are *non-degenerate* in the sense that the corresponding polynomial has d distinct roots. To make this precise, we will being by defining a subset of *degenerate problems* $\Sigma_d \subset \mathbb{C}^d$ and studying its properties. More precisely, Σ_d will be the set of polynomials which either have a repeated root or have degree not equal to d. As it turns out, Σ_d is has the structure of an affine algebraic variety—more precisely, a hypersurface: we'll define these terms later.

Recall that a univariate polynomial $f(x)$ has a repeated root if and only if $f(x)$ and its derivative $f'(x)$ share a common root. To understand for which polynomials this occurs, it will be useful to ask a more general question: given *two* univariate polynomials,

$$f = a_d x^d + \ldots + a_0, \tag{1}$$

$$g = b_e x^e + \ldots + b_0, \tag{2}$$

we ask: under what conditions on the coefficients $(a_d, \ldots, a_0, b_e, \ldots, b_0) \in \mathbb{C}^{d+e+2}$ do f and g share a common root? The answer can be given in terms of the classical $(d+e) \times (d+e)$ *Sylvester matrix*:

$$\operatorname{Syl}_{d,e}(f,g) = \begin{pmatrix} a_d & \cdots & a_1 & a_0 & 0 & \cdots & \cdots & 0 \\ & & & \ddots & & \ddots & & \\ 0 & \cdots & \cdots & 0 & a_d & \cdots & a_1 & a_0 \\ b_e & \cdots & b_0 & 0 & 0 & \cdots & \cdots & 0 \\ & & & \ddots & & \ddots & & \\ 0 & \cdots & & \cdots & 0 & b_e & \cdots & b_0 \end{pmatrix} \tag{3}$$

We define the *resultant* of f and g to be the determinant of the Sylvester matrix,

$$\operatorname{Res}_{d,e}(f,g) = \det \operatorname{Syl}_{d,e}(f,g). \tag{4}$$

Note that $\operatorname{Res}_{d,e}(f,g)$ may be viewed as a multivariate polynomial in the coefficients of f and g. When $g = f'$, cofactor expansion along the first column of $\operatorname{Syl}_{d,d-1}(f,f')$ gives us

$$\operatorname{Res}_{d,d-1}(f,f') = a_d \Delta_d(a_d, \ldots, a_0), \tag{5}$$

where $\Delta_d(a_0, \ldots, a_d)$ is a polynomial known as the *discriminant* of f.

Example 1. Letting $f = a_2 x^2 + a_1 x + a_0$, we have

$$\operatorname{Syl}_{2,1}(f,f') = \begin{pmatrix} a_2 & a_1 & a_0 \\ 2a_2 & a_1 & 0 \\ 0 & 2a_2 & a_1 \end{pmatrix}.$$

Cofactor expansion gives (at least up to sign) the familiar discriminant Δ_2:

$$\text{Res}_{2,1}(f, f') = a_2 \begin{vmatrix} a_1 & 0 \\ 2a_1 & a_1 \end{vmatrix} - 2a_2 \begin{vmatrix} a_1 & a_0 \\ 2a_2 & a_1 \end{vmatrix} = a_2 \underbrace{(4a_0a_2 - a_1^2)}_{\Delta_2}.$$

Definition 0.1. For $d \geq 1$, identifying $f = a_d x^d + \ldots + a_0$ with $(a_d, \ldots, a_0) \in \mathbb{C}^{d+1}$, we define

$$\Sigma_d = \{f \in \mathbb{C}^{d+1} \mid \text{Res}_{d,d-1}(f, f') = 0\} = \{f \in \mathbb{C}^{d+1} \mid \Delta_d(a_d, \ldots, a_0) = 0 \text{ or } a_d = 0\}.$$

Proposition 0.2. Two polynomials whose resultant is nonzero do not have a common root. In particular, if $f = a_d x^d + \ldots + a_0$ is a degree-d univariate polynomial with $\Delta_d(a_d, \ldots, a_0) \neq 0$, then f has no repeated roots.

Proof. Let f be as in 1 and g be as in 2, and suppose that $\text{Res}_{d,e}(f, g) \neq 0$. Then the linear system

$$\begin{bmatrix} \alpha_{e-1} & \cdots & \alpha_0 & \beta_{d-1} & \cdots & \beta_0 \end{bmatrix} \text{Syl}_{d,e}(f, g) = \begin{bmatrix} 0 & \cdots & 0 & 1 \end{bmatrix}$$

has a solution $\begin{bmatrix} \alpha_{e-1} & \cdots & \beta_0 \end{bmatrix}$. Observe that such a solution corresponds to a polynomial identity

$$(\alpha_{e-1}x^{e-1} + \ldots + \alpha_0)f(x) + (\beta_{d-1}x^{d-1} + \ldots + \beta_0)g(x) = 1. \tag{6}$$

If f and g had a common root, we could deduce $0 = 1$ from 6; therefore, no common root can exist. ☐

Homotopy continuation methods for computing roots of $g \in \mathbb{C}^{d+1}$ perform the following steps:

1. Pick some $f \in \mathbb{C}^{d+1} \setminus \Sigma_d$ whose roots we already know.
2. Set up a *homotopy function*

$$\begin{aligned} H : [0, 1] &\to \mathbb{C}^{d+1} \\ \text{such that } H(x; 0) = f(x),\ H(x; 1) &= g(x),\ H(x; t) \notin \Sigma_d\ \forall t \in [0, 1). \end{aligned} \tag{7}$$

3. For some discretization of the unit inverval

$$0 = t_0 < t_1 < \ldots < t_{k-1} < t_k = 1, \tag{8}$$

   use known roots of the equation $H(x; t_i) = 0$ to estimate the roots of $H(x; t_{i+1}) = 0$ for each $i = 0, \ldots, k - 1$.

The f and g in the homotopy function 7 are known as the *start system* and *target system*, respectively—we'll use the same terminology later when we construct homotopies for systems of multivariate equations.

Remark: Our use of the term "homotopy" is somewhat nonstandard: H is really a path in the space of polynomials $\mathbb{C}^{d+1}$. If we fix our start system $f \in \mathbb{C}^{d+1}$, the the homotopy function in 7 gives rise to a homotopy (in the traditional sense) on the space of polynomials $\mathbb{C}^{d+1}$, namely

$$\begin{aligned} H_f : \mathbb{C}^{d+1} \times [0, 1] &\to \mathbb{C}^{d+1} \\ (g, t) &\mapsto H(x; t). \end{aligned}$$

Note that 7 is only an abstract specification of a homotopy function; we haven't yet shown that such a function actually exists. Fortunately, there is a simple choice of start system that works for any $d \geq 1$:

$$f(x) = x^d - 1. \tag{9}$$

Observe that the d-th roots of unity

$$e^{2\pi ik/d} = \cos(2\pi ik/d) + i \sin(2\pi ik/d), \quad k = 0, \ldots, d - 1,$$

where $i^2 = -1$, are all roots of f. This follows from Euler's formula:

$$(e^{2\pi ik/d})^d - 1 = (e^{2\pi i})^k - 1 = 1^k - 1 = 0.$$

Furthermore, these are the only roots, due to the following general fact.

Proposition 0.3. Let $\mathbb{K}$ be a field. A univariate polynomial of degree d with coefficients has *at most* d roots in $\mathbb{K}$.

(This is easily proven using polynomial long division, which we will later generalize using Gröbner bases.) Our next observation is that the homotopy function in 7 can be understood as a *path* in the space of polynomials; our abstract specification requires that all points along this path except the target system don't lie in Σ_d. When the target system is also non-degenerate, we have the following connectivity result.

Proposition 0.4. Fix $f, g \in \mathbb{C}^{d+1} \setminus \Sigma_d$. Then there exists a path $H_{f,g} : [0,1] \to \mathbb{C}^{d+1} \setminus \Sigma_d$ with coordinate functions quadratic in t such that $H_{f,g}(0) = f$ and $H_{f,g}(1) = g$.

Proof. Consider the linear segment $s : [0,1] \to \mathbb{C}^{d+1}$ connecting f and g,

$$s(t) = s(x;t) = (1-t)f + tg, \tag{10}$$

and the following univariate polynomial in t:

$$h(t) = \mathrm{Res}_{d,d-1}\left(s(x;t), \frac{\partial}{\partial x} s(x;t)\right). \tag{11}$$

By Proposition 0.3, h has finitely-many complex roots: call them $t_j = x_j + iy_j$, for $j = 1, \ldots, m$. Consider the family of paths $\gamma_c : [0,1] \to \mathbb{C}$, parametrized by $c \in \mathbb{R}$, which are defined by $\gamma_c(t) = t + ict(1-t)$. We may then choose c so that $\gamma_c(t) \neq t_j$ for all $t \in [0,1]$ and $j = 1, \ldots, m$. Indeed, any

$$c < \min_{1 \le j \le m} \left(\frac{y_i}{x_i(1-x_i)}\right)$$

will work. Setting $H_{f,g}(t) = s(x; \gamma_c(t))$ then gives the result. □

Finally, before proving Theorem 1, we will need a rough bound on the size of the roots of a polynomial.

Proposition 0.5. [Cauchy Bound] For $f = a_d x^d + \ldots + a_0 \in \mathbb{C}[x]$ of degree d, any root x of f satisfies

$$|x| \le 1 + \max_{0 \le i \le d-1} \frac{|a_i|}{|a_d|}.$$

Proof of Theorem 1. Set $f(x) = x^d - 1$. We show that any $g \in \mathbb{C}^{d+1}$ has a root. To see this, consider first the case where $g \notin \Sigma_d$, and set $H(x;t) = H_{f(x),g(x)}(t)$, with $H_{f,g}$ as in the statement of Proposition 0.4.

Consider the set

$$D = \{t \in [0,1] \mid H(x;t) \text{ has a root } x \in \mathbb{C}\}. \tag{12}$$

To show that g has a root, it suffices to show that $D = [0,1]$. This will follow if we show that $D \subset [0,1]$ is a nonempty, open, and closed subset. Using Proposition 0.2, we have $0 \in D$, so D is nonempty. Furthermore, Proposition 0.2 implies that for any $t \in D$ we have

$$H(x;t) = 0 \quad \Rightarrow \quad \frac{\partial}{\partial x} H(x;t) \neq 0.$$

By the *implicit function theorem*, there exists an interval $I = (t - \epsilon, t + \epsilon) \subset \mathbb{R}$ such that $H(x;t')$ has a root for all $t' \in I$. This shows that D is open. Finally, to see that D is closed, consider the closed set

$$\mathcal{I} = \{(x,t) \in \mathbb{C} \times [0,1] \mid H(x;t) = 0\} \subset [0,1]. \tag{13}$$

The set $\mathcal{I}$ is also bounded; this follows from an application of Proposition 0.5 and the extreme value theorem. Now, if $t \in [0,1]$ is any limit point of the set D, there exists a sequence $(x_j, t_j)_{j=1}^{\infty} \in \mathcal{I}$ converging to a point $(t, x) \in \mathcal{I}$. Since $H(x;t) = 0$, this implies $t \in D$, and we conclude that D is closed.

Finally, we consider the case $g \in \Sigma_d$. We may assume that g has degree d (otherwise we are done by induction.) If we construct the homotopy $H(x;t)$ connecting f and g as before, then there are only finitely many points $t \in [0,1]$ for which $H(x;t) \notin \Sigma_d$. Restricting H to a suitably small closed subinterval of $[0,1]$ gives a homotopy satisfying the conditions of 7. A limiting argument similar to the previous one shows g has a complex root. □

Reference: Rojas, J. M. (2024). On the BCSS Proof of the Fundamental Theorem of Algebra. arXiv:2406.12198.