A Numerical Analyst's Tubular Neighborhood Theorem

Zhonggang Zeng *

October 10, 2016

Abstract

For complex analytic manifolds in Euclidean spaces, this paper establishes a version of Tubular Neighborhood Theorem that is specifically formulated for the applications in regularization of ill-posed algebraic problems for numerical computation. The theorem and the geometric analysis are indispensable for achieving accurate and stable numerical solutions of those highly sensitive problems using floating point arithmetic even if the data may be perturbed.

1 Introduction

When a scientific computing problem is attempted, there are at least three problems lurking around: The underlying problem that is intended but may be hidden, the given problem embodied in the available data, and the actual problem that ends up being solved. Those problems are not exactly the same but the differences are usually too small to be noticed until certain unexpected solutions emerge.

Think about a seemingly simple root-finding problem for a given polynomial

$$\tilde{p}(x) \quad = \quad x^5 + 4.667\, x^4 + 8.667\, x^3 + 8\, x^2 + 3.667\, x + .6667, \tag{1}$$

What is really intended may be roots of the hidden polynomial

$$p(x) \quad = \quad x^5 + \tfrac{14}{3}\, x^4 + \tfrac{26}{3}\, x^3 + 8\, x^2 + \tfrac{11}{3}\, x + \tfrac{2}{3} \quad \equiv \quad (x+1)^4 \left(x + \tfrac{2}{3}\right) \tag{2}$$

for which $\tilde{p}(x)$ serves as empirical data. Sending the data polynomial $\tilde{p}(x)$ to any commercial software such as Matlab for its roots, the results such as

$$-1.1211497 + 0.1032009i, \quad -1.1211497 - 0.1032009i, \quad -0.8863741 + 0.1451301i, \quad -0.8863741 - 0.1451301i, \quad -0.6519516$$

*Department of Mathematics, Northeastern Illinois University, Chicago, Illinois 60625, USA. email: zzeng@neiu.edu. Research is supported in part by NSF under grant DMS-1620337.

are far from expected even though they match exact roots of $\tilde{p}$ for the first three to four digits. The actual problem being solved is the exact root collection of the polynomial

$$\begin{aligned} \hat{p}(x) &= (x + 1.1211497 - 0.1032009i)(x + 1.1211497 + 0.1032009i) \cdots (x + 0.6519516) \\ &\approx x^5 + 4.666999\, x^4 + 8.666997\, x^3 + 7.999998\, x^2 + 3.6669996\, x + .6667001 \end{aligned}$$

All three problems are nearby as we expected, but the computing results bear no resemblance to the intended solutions $-1, \ -\frac{2}{3},$ as shown in Figure 1.

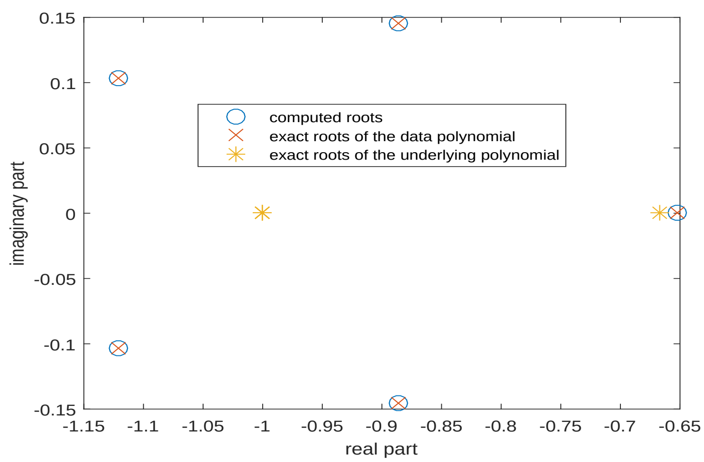

Figure 1: Computed roots and roots of the data polynomial $\tilde{p}$ are far from the roots of the underlying polynomial p due to hypersensitivity.

Dubbed as one of the "Maxims about numerical mathematics, science, computers and life on earth" [13]: "If the answer is highly sensitive to perturbations, you have probably asked the wrong question". Attempt in solving the root-finding problem of $\tilde{p}(x)$ in conventional sense is apparently such a wrong problem since a tiny perturbation from $p(x)$ to $\hat{p}(x)$ alters the solution seemingly completely.

There is also a fourth problem nearby. Wouldn't it make more sense to seek what we shall establish as the *regularized roots* of $\tilde{p}(x)$ and its *regularized factorization*

$$1.000068\,(x + 0.99994)^4\,(x + 0.66682),$$

particularly after proving such roots and factorization are existent, unique, Lipschitz continuous with respect to the data $\tilde{p}$, approximate those of the hidden polynomial $p(x)$, and can be computed accurately using floating point arithmetic from empirical data [15, 16]?

Those problems with extremely high sensitivities beyond finite bounds are known to be *ill-posed problems*, a concept attributed to Hadamard. As it turns out, such problems are abundant in scientific computing. Besides root-finding, even some of the most basic algebraic problems are ill-posed, such as matrix ranks and subspaces, polynomial greatest common divisors and factorizations, defective eigenvalues and Jordan Canonical Forms, etc. Those are the problems we will encounter inevitably in scientific computing.

Are the roots of aforementioned $p(x)$ really sensitive to data perturbations as alleged?

In a legendary technical report began circulating in 1972 [9] but never published, Kahan argues that it is a "misconception" to consider multiple roots hypersensitive to data perturbations, and he proves that the sensitivity of a root is finitely bounded if the data perturbation is constrained such that the multiplicity is preserved and thus the identity of the root remains intact. Kahan further argues that, albeit without proof, the ill-posed problems of a certain structure form what he called a "pejorative manifold". A problem is sensitive only if the data point is *near* the manifold but not necessarily so *on* it.

Hypersensitive problems frequently form complex analytic manifolds, such as those problems mentioned above. Furthermore, those manifolds entangle to form strata in which every manifold is embedded in the closure of some other manifolds of lower codimensions. As a result, the hypersensitivity of the problem is not random but directional: The singularity of a problem can only be decreased and never increases by a tiny perturbation. Using the root-finding example above, the roots of $p(x)$ changes in one and only one direction: The multiplicity 4 reduces by tiny perturbations.

The phenomenon that certain hypersensitive problems form a manifold, rather than merely an algebraic variety, is crucial in the analysis, regularization and numerical solution of those problems thanks to the existence of the *tubular neighborhood* that is considered "one of the most useful notions in the theory of differentiable manifolds" [5]. When the problem data are given as empirical and subject to small perturbations due to measurement and round-off, we are having a point near the complex analytic manifold in the problem space. Assuming the data are reasonably accurate so that the point remains in the tubular neighborhood of the manifold, the Tubular Neighborhood Theorem ensures the projection from the empirical data point to the manifold to be existent, unique and Lipschitz continuous, making the projection a well-posed least squares problem that is accurately solvable in numerical computation.

Geometric theories and insights have been applied in numerical analysis effectively in the works such as Sommese, Verschelde and Wampler[11], Corless, Galligo, Kotsireas and Watt [3], Edelman, Elmroth and Kågström [7, 8], Lippert and Edelman [10], Demmel and Edelman [4], Absil, Mahony and Sepulchre [1], Edelman, Arias and Smith [6], and so on. However, the tremendous advantage of the tubular neighborhoods has apparently not yet been utilized in numerical analysis. Specifically aimed at the application of solving hypersensitive algebraic problems in this paper, we establish a version of Tubular Neighborhood Theorem for complex analytic manifold in Euclidean spaces isometrically isomorphic to $\mathbb{C}^n$, and provide an elementary proof with a dose of flavor in numerical analysis. The theorem and the proof are not intended to impress differential topologists but to fill a gap in the regularization theory of hypersensitive problems and to complete the works of numerical factorization and numerical greatest common divisors of polynomials.

This paper attempts to provide a synergy between numerical analysis and differential topology. Besides establishing the version of Tubular Neighborhood Theorem for numerical analysis applications, we outline the geometric analysis of ill-posed problems in a series of observations, the regularization theory based on the geometry and the resulting strategy for solving the regularized problem accurately in numerical computation even if the data may be perturbed. These geometric analyses set the foundation for our software package NACLAB[1] for numerical algebraic computing.

2 Preliminaries

The space of n-dimensional vectors of complex numbers is denoted by $\mathbb{C}^n$ as a *topological vector space*. As a vector space, it is closed to addition and scalar multiplication. In the sense of a topological space, its topology is derived from the Euclidean norm $\|\cdot\|_2$.

We consider *finite-dimensional normed vector spaces* denoted by, say $\mathcal{V}$, $\mathcal{W}$ and so on. The norm $\|\mathbf{v}\|$ is understood from the context as *the* norm in the space where $\mathbf{v}$ belongs. Such a vector space $\mathcal{V}$ is isomorphic to $\mathbb{C}^n$ where $n = \dim(\mathcal{V})$, the dimension of $\mathcal{V}$.

A mapping $F : \Omega \subset \mathbb{C}^n \to \mathbb{C}^m$ is *holomorphic* if every one of the m components of F is a holomorphic function in every one of the n variables in the open subset Ω of $\mathbb{C}^n$. We can designate a variable name, say $\mathbf{z}$, for F and denote F as $\mathbf{z} \mapsto F(\mathbf{z})$. In that case the Jacobian matrix of F at any particular $\mathbf{z}_0 \in \Omega$ is denoted by $F_{\mathbf{z}}(\mathbf{z}_0)$.

Let $\mathcal{V}$ and $\mathcal{W}$ be finite-dimensional normed vector spaces isomorphic to $\mathbb{C}^n$ and $\mathbb{C}^m$ respectively via isomorphisms $\psi_{\mathcal{V}} : \mathcal{V} \to \mathbb{C}^n$ and $\psi_{\mathcal{W}} : \mathcal{W} \to \mathbb{C}^m$. Assume $\mathbf{g}$ is a mapping from an open subset Σ of $\mathcal{V}$ to $\mathcal{W}$ with a representation $\mathbf{z} \mapsto G(\mathbf{z})$ where $G : \psi_{\mathcal{V}}(\Sigma) \subset \mathbb{C}^n \to \mathbb{C}^m$ such that $\mathbf{g} = \psi_{\mathcal{W}}^{-1} \circ G \circ \psi_{\mathcal{V}}$. We say $\mathbf{g}$ is holomorphic in Σ if G is holomorphic in $\psi_{\mathcal{V}}(\Sigma)$. Denoting the variable of $\mathbf{g}$ as, say $\mathbf{v}$, the *Jacobian* of $\mathbf{g}$ at any particular $\mathbf{v}_0 \in \Sigma$ is defined as the linear transformation $\mathbf{g}_{\mathbf{v}}(\mathbf{v}_0) : \mathcal{V} \to \mathcal{W}$ in the form of

$$\mathbf{v} \longmapsto \mathbf{g}_{\mathbf{v}}(\mathbf{v}_0)(\mathbf{v}) \quad = \quad \psi_{\mathcal{W}}^{-1} \circ G_{\mathbf{z}}(\mathbf{z}_0) \circ \psi_{\mathcal{V}}(\mathbf{v})$$

where $\mathbf{z}_0 = \psi_{\mathcal{V}}(\mathbf{z}_0)$. The Jacobian $\mathbf{g}_{\mathbf{v}}(\mathbf{v}_0)$ so defined as a linear transformation is invariant under change of bases and the respective isomorphisms. Let $G_{\mathbf{z}}(\mathbf{z}_0)^{\mathsf{H}}$ and $G_{\mathbf{z}}(\mathbf{z}_0)^{\dagger}$ be the Hermitian transpose and the Moore-Panrose inverse of the Jacobian matrix $G_{\mathbf{z}}(\mathbf{z}_0)$ respectively. If we further assume the isomorphisms $\psi_{\mathcal{V}}$ and $\psi_{\mathcal{W}}$

[1]http://homepages.neiu.edu/∼naclab

are isometric so that $\|\psi_{\mathcal{V}}(\mathbf{v})\|_2 = \|\mathbf{v}\|$ and $\|\psi_{\mathcal{W}}(\mathbf{w})\|_2 = \|\mathbf{w}\|$ for all $\mathbf{v} \in \mathcal{V}$ and $\mathbf{w} \in \mathcal{W}$, then the Hermitian transpose and the Moore-Panrose inverse of $\mathbf{g}_{\mathbf{v}}(\mathbf{v}_0)$ are well-defined as linear transformations $\mathbf{g}_{\mathbf{v}}(\mathbf{v}_0)^{\mathsf{H}}$ and $\mathbf{g}_{\mathbf{v}}(\mathbf{v}_0)^{\dagger}$ respectively from $\mathcal{W}$ to $\mathcal{V}$ in the forms of

$$\begin{aligned} \mathbf{g}_{\mathbf{v}}(\mathbf{v}_0)^{\mathsf{H}} &= \psi_{\mathcal{V}}^{-1} \circ G_{\mathbf{z}}(\mathbf{z}_0)^{\mathsf{H}} \circ \psi_{\mathcal{W}} \\ \mathbf{g}_{\mathbf{v}}(\mathbf{v}_0)^{\dagger} &= \psi_{\mathcal{V}}^{-1} \circ G_{\mathbf{z}}(\mathbf{z}_0)^{\dagger} \circ \psi_{\mathcal{W}} \end{aligned}$$

that are invariant under any change of bases resulting in isometric isomorphisms. The holomorphic mapping $\mathbf{v} \mapsto \mathbf{g}(\mathbf{v})$ is said to be *non-degenerate* at $\mathbf{v}_0$ if the Jacobian $\mathbf{g}_{\mathbf{v}}(\mathbf{v}_0)$ is an injective linear transformation.

The notion of holomorphic mapping extends beyond open domains. A mapping $\mathbf{f} : \Pi \subset \mathcal{V} \to \mathcal{W}$ is holomorphic in a non-open domain Π if there is an open subset Ω of $\mathcal{V}$ containing Π and a holomorphic mapping $\mathbf{g}$ defined in Ω such that $\mathbf{f}(\mathbf{z}) \equiv \mathbf{g}(\mathbf{z})$ for all $\mathbf{z} \in \Pi$.

3 Complex analytic manifolds

In this paper, applications will be studied and computations will be conducted in finite-dimensional normed vector spaces in which we shall consider complex analytic manifolds in topological sense.

A finite-dimensional normed vector space $\mathcal{V}$ has an obvious topology in which an ε-neighborhood $N_\varepsilon(\mathbf{v}) := \{\mathbf{z} \in \mathcal{V} \mid \|\mathbf{z} - \mathbf{v}\| < \varepsilon\}$ of any $\mathbf{v} \in \mathcal{V}$ and $\varepsilon > 0$ is an open set along with unions of any collection of those neighborhoods. The vector space $\mathcal{V}$ equipped with this topology becomes a topological vector space. A subset $\Sigma \subset \mathcal{V}$ is a (topological) subspace of $\mathcal{V}$ with the subspace topology in which Λ is an open set if and only if there is an open set Ω in $\mathcal{V}$ such that $\Lambda = \Omega \cap \Sigma$.

The standard definition of manifolds in differential topology is more abstract and serves broader purposes than what we need for our applications in numerical solutions of ill-posed algebraic problems. We restrict our elaborations to the type of complex analytic manifolds in Euclidean spaces. They are, in fact, submanifolds of finite-dimensional normed vector spaces.

Definition 3.1 (Complex Analytic Manifold) *A topological subspace* Π *of a finite-dimensional normed vector space* $\mathcal{V}$ *over* $\mathbb{C}$ *is called a* complex analytic manifold *of dimension* m *if, for every point* $\mathbf{x} \in \Pi$, *there is an open neighborhood* Σ *of* $\mathbf{x}$ *in* $\mathcal{V}$ *and a holomorphic mapping* ϕ *from* $\Sigma \cap \Pi$ *into an open subset* Λ *of* $\mathbb{C}^m$ *with a holomorphic inverse* ψ *such that* $\phi \circ \psi(\mathbf{z}) \equiv \mathbf{z}$ *for all* $\mathbf{z} \in \Lambda$ *and* $\psi \circ \phi(\mathbf{y}) \equiv \mathbf{y}$ *for all* $\mathbf{y} \in \Sigma \cap \Pi$.

The dimension deficit $\dim(\mathcal{V}) - m$ is called the *codimension* of Π in $\mathcal{V}$ denoted by $\mathit{codim}(\Pi)$. The vector space $\mathcal{V}$ itself is a complex analytic manifold in $\mathcal{V}$ in which every complex analytic manifold is a submanifold.

Example 1 (Rank Manifold) [4] In the vector space $\mathbb{C}^{m\times n}$ of $m \times n$ matrices of complex entries with the Frobenius norm, the subset

$$\mathcal{C}_r^{m\times n} \;=\; \left\{ A \in \mathbb{C}^{m\times n} \;\middle|\; \mathit{rank}\,(\,A\,) = r \right\} \tag{3}$$

is a complex analytic manifold of codimension $(m-r)(n-r)$. Here $\mathit{rank}\,(\,A\,)$ denotes the rank of the matrix A. The proof is quite straightforward: For any matrix $A \in \mathcal{C}_r^{m\times n}$, there are r linearly independent columns forming a submatrix $X \in \mathbb{C}^{m\times r}$. The submatrix of the remaining $n-r$ columns can be written as $X\,Y$ where $Y \in \mathbb{C}^{r\times(n-r)}$. The mapping $A \mapsto (X, Y)$ is holomorphic and invertible from a neighborhood of A in Π into the vector space $C^{m\times r} \times \mathbb{C}^{r\times(n-r)}$ that is isometrically isomorphic to $\mathbb{C}^{mn-(m-r)(n-r)}$.

Example 2 (Univariate Factorization Manifold) [14] Let the vector space

$$\mathbb{P}_n \;=\; \left\{ a_0 + a_1\,x + \cdots + a_n\,x^n \;\middle|\; a_0, \ldots, a_n \in \mathbb{C} \right\}$$

with the norm $\|a_0 + a_1\,x + \cdots + a_n\,x^n\| \;=\; \sqrt{|a_0|^2 + \cdots + |a_n|^2}$, making $\mathbb{P}_n$ isometrically isomorphic to $\mathbb{C}^{n+1}$. For any positive integers $\ell_1 + \cdots + \ell_k = n$, the subset

$$\mathcal{F}_{\ell_1\cdots\ell_k} \;=\; \left\{ \alpha\,(x-z_1)^{\ell_1}\cdots(x-z_k)^{\ell_k} \in \mathbb{P}_n \;\middle|\; \alpha, z_1, \ldots, z_k \in \mathbb{C},\; z_i \neq z_j \text{ for } i \neq j \right\} \tag{4}$$

is a complex analytic manifold of codimension $n-k$ in $\mathbb{P}_n$.

Example 3 (Univariate GCD Manifold) [18] The product space $\mathbb{P}_m \times \mathbb{P}_n$ with the norm $\big\|(p,q)\big\| \;=\; \sqrt{\|p\|^2 + \|q\|^2}$ is isometrically isomorphic to $\mathbb{C}^{m+n+2}$. For every pair $(p,q) \in \mathbb{P}_m \times \mathbb{P}_n$, let $\mathit{gcd}\,(p,q)$ denote the greatest common divisor of p and q. The subset

$$\mathscr{P}_{m,n}^{k} \;=\; \left\{ (p,q) \in \mathbb{P}_m \times \mathbb{P}_n \;\middle|\; \mathit{deg}(p) = m,\; \mathit{deg}(q) = n,\; \mathit{deg}(\mathit{gcd}\,(p,q)) = k \right\}$$

is a complex analytic manifold of codimension k in $\mathbb{P}_m \times \mathbb{P}_n$ where $\mathit{deg}(\cdot)$ is the degree of any polynomial $(\cdot)$.

Not all algebraic varieties are manifolds, as shown in the following example.

Example 4 (Example of a non-manifold) Consider the subset

$$\Pi_5^2 \;=\; \big\{(x-\alpha)^2\, q(x) \in \mathbb{P}_5 \;\big|\; q \in \mathbb{P}_3,\; q(\alpha) \neq 0\big\}$$

of $\mathbb{P}_5$ having a double root and define the mapping

$$\mathbf{f} \;:\; (\alpha, a_0, a_1, a_2, a_3) \;\longmapsto\; (x-\alpha)^2\,(a_0 + a_1\,x + a_2\,x^2 + a_3\,x^3)$$

from $\mathbb{C}^5$ to $\mathbb{P}_5$. The polynomial

$$\begin{aligned} p(x) &= (x-1)^2\,(x-2)^2\,(x-3) \\ &\equiv (x-1)^2\,(-12 + 16\,x - 7\,x^2 + x^3) \\ &\equiv (x-2)^2\,(-3 + 7\,x - 5\,x^2 + x^3) \end{aligned}$$

in Π_5^2 is the image of both $\mathbf{f}(1,-12,16,-7,1)$ and $\mathbf{f}(2,-3,7,-5,1)$. Namely, the polynomial p is in the intersection $\mathbf{f}(\Omega_1) \cap \mathbf{f}(\Omega_2)$ of two branches of Π_5^2 where $\Omega_1, \Omega_2 \subset \mathbb{C}^5$ are small neighborhoods of $(1,-12,16,-7,1)$ and $(2,-3,7,-5,1)$ respectively. As a result, the subset Π_5^2 is not a complex analytic manifold.

4 The Tubular Neighborhood Theorem

The Tubular Neighborhood Theorem can be illustrated in Figure 2: A complex analytic manifold is contained in an open *tubular neighborhood* in which every point can be uniquely projected onto the manifold following a normal line. For our application of numerical solution of ill-posed problems, it is sufficient for the projection from the tubular neighborhood to the manifold to be Lipschitz continuous. Being smooth or analytic is a unnecessary luxury. The Lipschitz constant can also serve as a condition number.

The Tubular Neighborhood Theorem is one of the fundamental results in differential topology. Standard versions of the Tubular Neighborhood Theorem for real smooth manifolds can be found in textbooks of differential geometry. Those versions are presented in abstract forms for general purposes and do not appear directly applicable to our applications involving complex analytic manifolds. The following Theorem 4.1 is a down-to-earth version of the Tubular Neighborhood Theorem for submanifolds of normed vector spaces isometrically isomorphic to $\mathbb{C}^n$'s and it is specifically tailored for our applications. Since we can't find such a version in the literature, we provide a numerical analysis proof using the Gauss-Newton iteration that also serves as the mechanism for calculating the projection to the manifold.

Theorem 4.1 (Tubular Neighborhood Theorem) *Let* Π *be a complex analytic manifold in a normed vector space* $\mathcal{V}$ *that is isometrically isomorphic to* $\mathbb{C}^n$. *There*

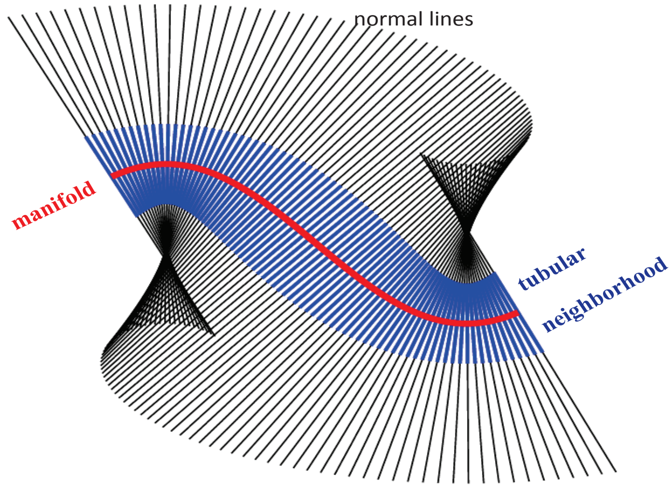

Figure 2: Illustration of the tubular neighborhood

is a tubular neighborhood*, namely an open subset* Ω *of* $\mathcal{V}$ *such that* $\Pi \subset \Omega$*, and every* $\mathbf{b} \in \Omega$ *has a unique projection* $\mathbf{x_b} \in \Pi$ *of minimum distance to* Π*, that is*

$$\left\|\mathbf{x_b} - \mathbf{b}\right\| \quad = \quad \inf_{\mathbf{x} \in \Pi} \left\|\mathbf{x} - \mathbf{b}\right\|. \tag{5}$$

Furthermore, the projection $\mathbf{b} \mapsto \mathbf{x_b}$ *from* Ω *to* Π *is locally Lipschitz continuous.*

We shall prove this particular version of Tubular Neighborhood Theorem in §6.

From the perspective of a numerical analyst, the projection from any point $\mathbf{b}$ in the tubular neighborhood Ω to the manifold Π is the (nonlinear) least squares problem of solving the (nonlinear) system $\mathbf{f}(\mathbf{x}) = \mathbf{b}$ where $\mathbf{f} : \mathbb{C}^m \to \mathcal{V}$ is the local diffeomorphism for the manifold Π. The Tubular Neighborhood Theorem ensures the solution of such a least squares problem exists, is unique and is Lipschitz continuous.

5 The nonlinear least squares problem

Let $\mathcal{V}$ and $\mathcal{W}$ be normed vector spaces over $\mathbb{C}$ that are isometrically isomorphic to $\mathbb{C}^n$ and $\mathbb{C}^m$ respectively with $m > n$. For any holomorphic mapping $\mathbf{x} \mapsto \mathbf{f}(\mathbf{x})$ from an open subset Ω of $\mathcal{V}$ to $\mathcal{W}$, consider the equation

$$\mathbf{f}(\mathbf{x}) \quad = \quad \mathbf{b}. \tag{6}$$

Since the image $\mathbf{f}(\Omega)$ is some kind of surface of dimension at most n in the m-dimensional space $\mathcal{W}$ with $m > n$, conventional solutions do not exist for the equation (6) except in special cases where $\mathbf{b} \in \mathbf{f}(\Omega)$. Instead, we seek a *least squares solution* $\mathbf{x}_* \in \Omega$ such that

$$\left\|\mathbf{f}(\mathbf{x}_*) - \mathbf{b}\right\|^2 \;=\; \min_{\mathbf{x}\in\Omega} \left\|\mathbf{f}(\mathbf{x}) - \mathbf{b}\right\|^2.$$

In other words, we seek $\mathbf{x}_*$ so that $\mathbf{f}(\mathbf{x}_*)$ is the projection of $\mathbf{b}$ on to the surface $\mathbf{f}(\Omega)$ so that the minimal distance from $\mathbf{b}$ to the surface is attained.

It is a straightforward verification that such a least squares solution is a *critical point* for the equation (6), namely

$$\mathbf{f}_{\mathbf{x}}(\mathbf{x}_*)^{\mathsf{H}} \big(\mathbf{f}(\mathbf{x}_*) - \mathbf{b}\big) \;=\; \mathbf{0}. \tag{7}$$

The interpretation of (7) is quite clear: At the least squares solution $\mathbf{x}_*$, the vector $\mathbf{f}(\mathbf{x}_*) - \mathbf{b}$ must be orthogonal to the tangent plane of the surface $\mathbf{f}(\Omega)$ at the point $\mathbf{f}(\mathbf{x}_*)$.

An effective method for finding the least squares solution of (6) is the *Gauss-Newton iteration*

$$\mathbf{x}_{k+1} \;=\; \mathbf{x}_k - \mathbf{f}_{\mathbf{x}}(\mathbf{x}_k)^\dagger \big(\mathbf{f}(\mathbf{x}_k) - \mathbf{b}\big) \qquad \text{for} \quad k = 0, 1, \ldots \tag{8}$$

The Gauss-Newton iteration (8) is locally convergent requiring the initial iterate $\mathbf{x}_0$ to be sufficiently near the least squares solution $\mathbf{x}_*$ *and* the residual $\left\|\mathbf{f}(\mathbf{x}_*) - \mathbf{b}\right\|$ to be sufficiently small. The following lemma provides detailed convergence conditions in Kantorovich style.

Lemma 5.1 [17] *Let $\mathcal{V}$ and $\mathcal{W}$ be finite-dimensional normed vector spaces isometrically isomorphic to $\mathbb{C}^n$ and $\mathbb{C}^m$ respectively. Assume $\mathbf{x} \mapsto \mathbf{f}(\mathbf{x})$ is a holomorphic mapping from an open subset $\Omega \subset \mathcal{V}$ to $\mathcal{W}$ with a critical point $\mathbf{x}_* \in \Omega$ of the system $\mathbf{f}(\mathbf{x}) = \mathbf{b}$. Assume $\Lambda \subset \Omega$ is an open neighborhood of $\mathbf{x}_*$ so that $\mathbf{f}$ is non-degenerate in Λ and there exist constants $\zeta, \gamma > 0$ such that*

$$\left\|\mathbf{f}_{\mathbf{x}}(\mathbf{z})^\dagger\right\| \;\le\; \zeta, \tag{9}$$

$$\left\|\mathbf{f}(\mathbf{z}) - \mathbf{f}(\tilde{\mathbf{z}}) - \mathbf{f}_{\mathbf{x}}(\tilde{\mathbf{z}})(\mathbf{z} - \tilde{\mathbf{z}})\right\| \;\le\; \gamma \left\|\mathbf{z} - \tilde{\mathbf{z}}\right\|^2 \tag{10}$$

for all $\mathbf{z}, \tilde{\mathbf{z}} \in \Lambda$. Further assume $\left\|\mathbf{f}(\mathbf{x}_*) - \mathbf{b}\right\|$ is small so that

$$\left\|\big(\mathbf{f}_{\mathbf{x}}(\mathbf{z})^\dagger - \mathbf{f}_{\mathbf{x}}(\mathbf{x}_*)^\dagger\big)\big(\mathbf{f}(\mathbf{x}_*) - \mathbf{b}\big)\right\| \;\le\; \sigma \left\|\mathbf{z} - \mathbf{x}_*\right\| \tag{11}$$

for a constant $\sigma < 1$ at every $\mathbf{z} \in \Lambda$. Then for all $\mathbf{x}_0 \in \Lambda$ satisfying

$$\left\|\mathbf{x}_0 - \mathbf{x}_*\right\| \;<\; \tfrac{1-\sigma}{\zeta\gamma} \tag{12}$$

and $\left\{\mathbf{x} \in \mathcal{V} \;\middle|\; \left\|\mathbf{x} - \mathbf{x}_*\right\| < \left\|\mathbf{x}_0 - \mathbf{x}_*\right\|\right\} \subset \Lambda$, the Gauss-Newton iteration (8) is well defined in Λ, converges to $\mathbf{x}_*$, and satisfies

$$\left\|\mathbf{x}_{k+1} - \mathbf{x}_*\right\| \;\le\; \big(\sigma + \zeta\gamma \left\|\mathbf{x}_k - \mathbf{x}_*\right\|\big) \left\|\mathbf{x}_k - \mathbf{x}_*\right\|$$

for $k = 0, 1, \ldots$ with $\mu = \sigma + \zeta\gamma \left\|\mathbf{x}_0 - \mathbf{x}_*\right\| < 1$.

If it is convergent, the Gauss-Newton iteration converges at a linear rate in general and at a quadratic rate if the residual $\|\mathbf{f}(\mathbf{x}_*) - \mathbf{b}\| = 0$ in the cases where the least squares solution is a conventional solution.

6 An elementary proof of the Tubular Neighborhood Theorem

Proof of Theorem 4.1. Let $\mathbf{z}_0$ be any particular point in Π. Since Π is a complex analytic manifold in $\mathcal{V}$, there is a bounded neighborhood $\mathcal{M}$ of $\mathbf{z}_0$ in $\mathcal{V}$, an open subset $\mathcal{N}$ of $\mathbb{C}^m$ where $m = \dim(\Pi)$, and a holomorphic mapping $\mathbf{x} \mapsto \mathbf{f}(\mathbf{x})$ from $\mathcal{N}$ to $\mathcal{M} \cap \Pi$ with a holomorphic inverse $\mathbf{f}^{-1}$. Denote $\mathbf{f}^{-1}(\mathbf{z}_0) = \mathbf{x}_0$. We can assume $\mathcal{M}$ is sufficiently small and there are constants $\zeta, \gamma > 0$ such that

$$\begin{aligned} \left\|\mathbf{f}_\mathbf{x}(\mathbf{y})^\dagger\right\| &\leq \zeta \\ \left\|\mathbf{f}(\mathbf{y}) - \mathbf{f}(\tilde{\mathbf{y}}) - \mathbf{f}_\mathbf{x}(\tilde{\mathbf{y}})(\mathbf{y} - \tilde{\mathbf{y}})\right\| &\leq \gamma \|\mathbf{y} - \tilde{\mathbf{y}}\|^2 \end{aligned}$$

for all $\mathbf{y}, \tilde{\mathbf{y}} \in \mathcal{N}$. Let $N_\varepsilon(\mathbf{x}_0) = \left\{\mathbf{x} \in \mathbb{C}^m \;\middle|\; \|\mathbf{x} - \mathbf{x}_0\| < \varepsilon\right\}$ and assume $\varepsilon > 0$ to be sufficiently small so that $N_\varepsilon(\mathbf{x}_0) \subset \mathcal{N}$ and $\|\mathbf{y} - \tilde{\mathbf{y}}\| < \frac{1}{2\zeta\gamma}$ for any $\mathbf{y}, \tilde{\mathbf{y}} \in N_\varepsilon(\mathbf{x}_0)$. Set $\varepsilon' = \frac{1}{3}\varepsilon$.

Let $\tilde{\mathcal{M}}$ be an open subset of $\mathcal{M}$ with $\mathbf{z}_0 \in \tilde{\mathcal{M}}$ and $\mathbf{f}^{-1}(\tilde{\mathcal{M}} \cap \Pi) \subset N_{\varepsilon'}(\mathbf{x}_0)$. Then there is an open subset $\hat{\mathcal{M}}$ of $\tilde{\mathcal{M}}$ with $\mathbf{z}_0 \in \hat{\mathcal{M}}$ such that, for every $\mathbf{b} \in \hat{\mathcal{M}}$,

$$\|\mathbf{b} - \mathbf{z}_0\| < \|\mathbf{b} - \mathbf{d}\| \qquad \text{for every} \qquad \mathbf{d} \in \mathcal{V} \setminus \tilde{\mathcal{M}}. \tag{13}$$

Since $\overline{\tilde{\mathcal{M}}} \cap \Pi$ is compact, there is a $\mathbf{z}_\mathbf{b} \in \overline{\tilde{\mathcal{M}}} \cap \Pi$ such that

$$\|\mathbf{b} - \mathbf{z}_\mathbf{b}\| = \min_{\mathbf{u} \in \overline{\tilde{\mathcal{M}}} \cap \Pi} \|\mathbf{b} - \mathbf{u}\| \leq \|\mathbf{b} - \mathbf{z}_0\|$$

and thus $\mathbf{z}_\mathbf{b} \in \tilde{\mathcal{M}} \cap \Pi$ due to (13). We have thus proved that, for every $\mathbf{b} \in \hat{\mathcal{M}}$, there exists an $\mathbf{x}_\mathbf{b} \in N_{\varepsilon'}(\mathbf{x}_0)$ such that $\mathbf{f}(\mathbf{x}_\mathbf{b}) = \mathbf{z}_\mathbf{b}$ and

$$\|\mathbf{f}(\mathbf{x}_\mathbf{b}) - \mathbf{b}\| = \min_{\mathbf{u} \in \Pi} \|\mathbf{b} - \mathbf{u}\| = \min_{\mathbf{x} \in \mathcal{N}} \|\mathbf{f}(\mathbf{x}) - \mathbf{b}\|$$

and such an $\mathbf{x}_\mathbf{b}$ can only occur in $N_{\varepsilon'}(\mathbf{x}_0)$. In other words, the least squares solution to the equation $\mathbf{f}(\mathbf{x}) = \mathbf{b}$ and $\mathbf{z}_\mathbf{b} = \mathbf{f}(\mathbf{x}_\mathbf{b})$ exists for every $\mathbf{b} \in \hat{\mathcal{M}}$ and $\mathbf{x}_\mathbf{b} \in N_{\varepsilon'}(\mathbf{x}_0)$.

We can further assume $\tilde{\mathcal{M}}$ to be sufficiently small so that, for all $\mathbf{b} \in \hat{\mathcal{M}} \subset \tilde{\mathcal{M}}$,

$$\left\|\left(\mathbf{f}_\mathbf{x}(\mathbf{y})^\dagger - \mathbf{f}_\mathbf{x}(\mathbf{z})^\dagger\right)\left(\mathbf{f}(\mathbf{z}) - \mathbf{b}\right)\right\| \leq \frac{1}{2}\|\mathbf{y} - \mathbf{z}\|$$

whenever $\mathbf{y} \in N_\varepsilon(\mathbf{x}_0)$ and $\mathbf{z} \in N_{\varepsilon'}(\mathbf{x}_0)$, making

$$\|\mathbf{y}_0 - \mathbf{x_b}\| \;<\; \frac{1}{2\zeta\gamma} \;=\; \frac{1-\frac{1}{2}}{\zeta\gamma}$$

for every $\mathbf{y}_0 \in N_{\varepsilon'}(\mathbf{x}_0)$. Furthermore, the inequality $\|\mathbf{x} - \mathbf{x_b}\| < \|\mathbf{y}_0 - \mathbf{x_b}\|$ with $\mathbf{y}_0 \in N_{\varepsilon'}(\mathbf{x}_0)$ implies

$$\begin{aligned}
\|\mathbf{x} - \mathbf{x}_0\| &\leq \|\mathbf{x} - \mathbf{x_b}\| + \|\mathbf{x_b} - \mathbf{x}_0\| \\
&< \|\mathbf{y}_0 - \mathbf{x_b}\| + \|\mathbf{x_b} - \mathbf{x}_0\| \\
&\leq \|\mathbf{y}_0 - \mathbf{x}_0\| + \|\mathbf{x}_0 - \mathbf{x_b}\| + \varepsilon' \\
&< \varepsilon' + \varepsilon' + \varepsilon' \;=\; \varepsilon.
\end{aligned}$$

Namely $\big\{\mathbf{x} \in \mathcal{V} \,\big|\, \|\mathbf{x}-\mathbf{x_b}\| < \|\mathbf{y}_0-\mathbf{x_b}\|\big\} \subset N_\varepsilon(\mathbf{x}_0)$. As a result, the Gauss-Newton iteration

$$\mathbf{y}_{k+1} \;=\; \mathbf{y}_k - \mathbf{f_x}(\mathbf{y}_k)^\dagger\left(\mathbf{f}(\mathbf{y}_k) - \mathbf{b}\right) \quad \text{for} \quad k = 0, 1, \ldots \tag{14}$$

converges to $\mathbf{x_b}$ from every $\mathbf{y}_0 \in N_{\varepsilon'}(\mathbf{x}_0)$. This least squares solution $\mathbf{x_b}$ is unique for every $\mathbf{b} \in \hat{\mathcal{M}}$. In fact, another least squares solution $\tilde{\mathbf{x}}_\mathbf{b}$ to the equation $\mathbf{f}(\mathbf{x}) = \mathbf{b}$ implies the Gauss-Newton iteration (14) converges to $\mathbf{x_b}$ from the initial iterate $\mathbf{y}_0 = \tilde{\mathbf{x}}_\mathbf{b}$. On the other hand, the Gauss-Newton iteration starting from the critical point $\tilde{\mathbf{x}}_\mathbf{b}$ stays at $\tilde{\mathbf{x}}_\mathbf{b}$, leading to $\tilde{\mathbf{x}}_\mathbf{b} = \mathbf{x_b}$.

We have now established a mapping $\mathbf{b} \mapsto \mathbf{x_b}$ from the domain $\hat{\mathcal{M}}$ to the manifold Π. We claim this mapping is Lipschitz continuous. Let $\mathbf{x_d} \in N_{\varepsilon'}(\mathbf{x}_0)$ be the least squares solution to the equation $\mathbf{f}(\mathbf{x}) = \mathbf{d}$ and perform one step of the Gauss-Newton iteration

$$\mathbf{x}_1 \;=\; \mathbf{x_b} - \mathbf{f_x}(\mathbf{x_b})^\dagger\big(\mathbf{f}(\mathbf{x_b}) - \mathbf{d}\big)$$

for the equation $\mathbf{f}(\mathbf{x}) = \mathbf{d}$ from the initial iterate $\mathbf{x_b}$. Using the identity $\mathbf{x_b} = \mathbf{x_b} - \mathbf{f_x}(\mathbf{x_b})^\dagger\big(\mathbf{f}(\mathbf{x_b}) - \mathbf{b}\big)$, we have

$$\|\mathbf{x}_1 - \mathbf{x_b}\| \;\leq\; \big\|\mathbf{f_x}(\mathbf{x_b})^\dagger\big\|\, \|\mathbf{b} - \mathbf{d}\|.$$

On the other and, Lemma 5.1 implies that there is a $\mu \leq \frac{1}{2} + \zeta\gamma\|\mathbf{x_b} - \mathbf{x_d}\|$ such that

$$\|\mathbf{x}_1 - \mathbf{x_d}\| \;\leq\; \mu\,\|\mathbf{x_b} - \mathbf{x_d}\|.$$

Consequently

$$\begin{aligned}
\|\mathbf{x_d} - \mathbf{x_b}\| &\leq \|\mathbf{x_b} - \mathbf{x}_1\| + \|\mathbf{x}_1 - \mathbf{x_d}\| \\
&\leq \big\|\mathbf{f_x}(\mathbf{x_b})^\dagger\big\|\, \|\mathbf{b} - \mathbf{d}\| + \mu\,\|\mathbf{x_b} - \mathbf{x_d}\|
\end{aligned}$$

and thus

$$\|\mathbf{x_b} - \mathbf{x_d}\| \;\leq\; \frac{1}{1-\mu}\,\big\|\mathbf{f_x}(\mathbf{x_b})^\dagger\big\|\, \|\mathbf{b} - \mathbf{d}\|. \tag{15}$$

The assymptotic upper bound of μ is $\frac{1}{2}$ when $\mathbf{d}$ approaches to $\mathbf{b}$ and μ can be arbitrarily small if $\mathbf{b}$ is close to the manifold Π.

Let Ω be the tubular neighborhood as the union of all those $\hat{\mathcal{M}}$ over all $\mathbf{z}_0$. The theorem is thus proved. $\square$

7 Geometry of hypersensitive problems

Attributed to Hadamard, the notion of *well-posed problem* requires a mathematical model to have a solution that satisfies existence, uniqueness and some kind of continuity with respect to data. For the problem to be feasible in numerical computation, its solution needs to be Lipschitz continuous so that the sensitivity is bounded with respect to data perturbation and round-off. The version of Tubular Neighborhood Theorem in Theorem 4.1 is specifically formulated and established to ensure that the (numerical) least squares problem of computing the nearest point on a complex analytic manifold from any point in the tubular neighborhood is a well-posed problem.

A problem whose solution lacks one of the existence, uniqueness or continuity is called an *ill-posed problem*. In particular, if the solution is not Lipschitz continuous with respect to data, we say the problem is *hypersensitive* with a condition number infinity. A typical textbook example in numerical analysis is the root-finding problem of polynomials in the form of solving the equation

$$a_0 + a_1 z + a_2 z^2 + \cdots + a_n z^n \quad = \quad 0 \qquad (16)$$

with $a_n \neq 0$. If the roots of the polynomial $p(z) = a_0 + a_1 z + a_2 z^2 + \cdots + a_n z^n$ are all simple, every solution z_* of (16) has a finite condition number $\frac{1}{|p'(z_*)|}$ and is thus Lipschitz continuous with respect the coefficient vector $(a_0, a_1, \ldots, a_n)$ serving as data. When multiple roots are present, however, the Lipschitz continuity is lost and the root-finding problem (16) becomes hypersensitive.

Geometrically, every such a polynomial p is a point in the vector space $\mathbb{P}_n$ that consists of polynomials with degrees up to n. If p has k roots of multiplicities $\ell_1, \ldots, \ell_k$ respectively where $\ell_1 + \cdots + \ell_k = n$, then p is on the complex analytic manifolds $\mathcal{F}_{\ell_1 \cdots \ell_k}$ defined in (4) with codimension $n-k$. The root-finding problem for $p(z) = 0$ is hypersensitive except for p ton be on the trivial factorization manifold $\mathcal{F}_{1\,1 \cdots 1}$, i.e. $k = n$, $\ell_1 = \cdots = \ell_n = 1$ and $\mathit{codim}(\mathcal{F}_{\ell_1 \cdots \ell_n}) = n - n = 0$. We can quantify the hypersensitivity of a polynomial p in terms of root-finding by defining the *singularity* of $p \in \mathcal{F}_{\ell_1 \cdots \ell_n})$ as $\mathit{codim}(\mathcal{F}_{\ell_1 \cdots \ell_n})$. We also say the *singularity* of the complex analytic manifold $\mathcal{F}_{\ell_1 \cdots \ell_n}$ is $n-k$. Notice that the trivial manifold $\mathcal{F}_{1\,1 \cdots 1}$ of singularity zero is open and dense in $\mathbb{P}_n$, leading to the observation:

> *Nonsingular problems are open and dense in the problem space. A problem is hypersensitive only if the data point resides on a manifold of nonzero singularity (codimension).*

When a simple root z_* of p is ill-conditioned, namely $p'(z_*) \approx 0$. The polynomial p is clearly near a polynomial with multiple roots and is thus near a manifold $\mathcal{F}_{\ell_1 \cdots \ell_n}$ of nonzero singularity $n-k$, implying:

> *A problem is ill-conditioned only if the data point is near a manifold of nonzero singularity (codimension).*

Consider the holomorphic mapping

$$\begin{array}{rccl} \mathbf{f} & : & \mathbb{C}^{k+1} & \longrightarrow & \mathbb{P}_n \\ & & (z_0, z_1, \ldots, z_k) & \longmapsto & z_0(x-z_1)^{\ell_1}\cdots(x-z_k)^{\ell_k}. \end{array}$$

For every $p \in \mathcal{F}_{\ell_1\cdots\ell_k}$, there is a $\mathbf{z} \in \mathbb{C}^{k+1}$ such that $\mathbf{f}(\mathbf{z}) = p$ and there are neighborhoods Λ and Δ of p and $\mathbf{z}$, respectively, along with a holomorphic mapping $\mathbf{g} : \Lambda \to \Delta$ that is the inverse of $\mathbf{f}$ in Δ. As a result, the mapping $\mathbf{f}$ serves as the local diffeomorphism for the complex analytic manifold $\mathcal{F}_{\ell_1\cdots\ell_k}$. By the proof of the Tubular Neighborhood Theorem, we have

$$\|\mathbf{z}-\tilde{\mathbf{z}}\|_2 \;\le\; \frac{1}{1-\mu}\left\|\mathbf{f}_{\mathbf{z}}(\mathbf{z})^{\dagger}\right\|\,\|p-\tilde{p}\|$$

for $\tilde{p} = \mathbf{f}(\tilde{\mathbf{z}})$ near p and $\mu \to 0$ when $\tilde{p}$ approaches p on the manifold $\mathcal{F}_{\ell_1\cdots\ell_k}$. Namely, even though the roots of the polynomial p are hypersensitive to *arbitrary perturbations* on p, they are in the solution of the equation $\mathbf{f}(\mathbf{z}) = p$ that is Lipschitz continuous for p constrained on the manifold $\mathcal{F}_{\ell_1\cdots\ell_k}$, validating Kahan's observation in 1972 [9]:

> *A singular problem is not hypersensitive if the perturbation is constrained such that the data point stays on its native manifold.*

Perhaps the most crucial observation of hypersensitive problems is the geometric explanation of the hypersensitivity: A problem is hypersensitive because it resides on a complex analytic manifold of nonzero singularity (codimension). Due to the dimension deficit, an arbitrary perturbation generically pushes the data (point) away from the native manifold and thus changes the solution structure completely no matter how tiny the perturbation is. Looking deeper, however, we need to point out that the solution changes from tiny perturbations in one and only one direction in reducing the singularity. Consider a polynomial, say

$$p(x) \;=\; (x-1)^2(x-2)^3 \;\in\; \mathcal{F}_{2\,3}$$

of singularity $5-2=3$. It is arbitrarily close to the infinitely many polynomials of different multiplicity structures such as

$$p_\varepsilon(x) \;=\; (x-1-\varepsilon)(x-1+\varepsilon)(x-2)^3 \;\in\; \mathcal{F}_{1\,1\,3}$$

of singularity $5-3=2$ when $\varepsilon \to 0$. Clearly the manifold $\mathcal{F}_{2\,3} \subset \overline{\mathcal{F}_{1\,1\,3}}$. We say a manifold Π is *embedded* in the manifold Σ, denoted by $\Pi \hookrightarrow \Sigma$ if the former is in the closure of the latter. It is common for hypersensitive problems to form complex analytic manifolds that entangle and form *strata* in which a singular manifold is embedded in manifolds of lower singularities. Figure 3 shows the stratification of manifolds of degree-5 polynomials.

As illustrated in Figure 3, the manifold $\mathcal{F}_{2\,3}$ is embedded in $\mathcal{F}_{1\,1\,3}, \mathcal{F}_{1\,2\,2}, \mathcal{F}_{1\,1\,1\,2}$ and the nonsingular $\mathcal{F}_{1\,1\,1\,1\,1}$. As a result, a small perturbation Δp on $p \in \mathcal{F}_{2\,3}$ can push the polynomial p to $p+\Delta p$ into those manifolds of lower singularities. There is a minimum distance δ from p to $\mathcal{F}_{1\,4}$ and $\mathcal{F}_5$ of equal or higher singularities and the perturbation Δp must have a magnitude at least δ to change the multiplicity structure of p to equal or higher singularities. Thus:

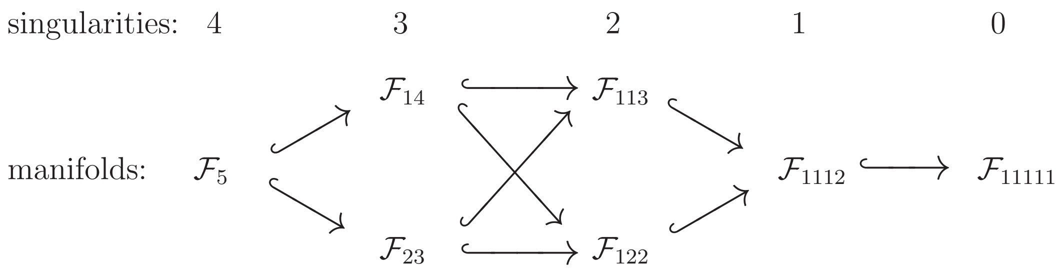

Figure 3: Stratification of manifolds of degree-5 polynomials and their singularities.

> *A small perturbation on the data of a hypersensitive problem can only reduce the singularity and never increase it.*

This unidirectional perturbation property of singularity is the most crucial property for numerical computation and provides a mechanism for identifying the manifold on which the problem resides. We define the distance of a point $\mathbf{z}$ from a manifold Π in a normed vector space as

$$\mathit{dist}\,(\,\mathbf{z},\,\Pi\,) \quad := \quad \inf_{\mathbf{y}\in\Pi} \|\mathbf{y}-\mathbf{z}\|.$$

Such a distance is zero if and only if $\mathbf{z} \in \overline{\Pi}$. For instance, a polynomial $p \in \mathcal{F}_{23}$ means that p has a zero distance to manifolds $\mathcal{F}_{23}$, $\mathcal{F}_{113}$, $\mathcal{F}_{122}$, $\mathcal{F}_{1112}$ and $\mathcal{F}_{11111}$. Among those manifolds, the native manifold $\mathcal{F}_{23}$ distinguishes itself as the one with the highest singularity

$$\mathit{codim}\big(\mathcal{F}_{23}\big) \quad = \quad \max_{\mathit{dist}\big(\,p,\,\mathcal{F}_{\ell_1\cdots\ell_k}\,\big)=0,\ \ell_1+\cdots+\ell_k=5} \mathit{codim}\big(\mathcal{F}_{\ell_1\cdots\ell_k}\big)$$

leading to the observation:

> *The native manifold of a hypersensitive problem is identifiable as the manifold with the highest singularity among all manifolds with a zero distance to the problem.*

In practical computation, a hypersensitive problem is likely represented by imperfect empirical data while the precise values are unknown. Even if the exact data are known, numerical algorithms inevitably introduce round-off errors during the solution process. Continuing the example of the root-finding problem, even if the polynomial $p \in \mathcal{F}_{23}$ is represented by the empirical data point $p+\Delta p$ where Δp is arbitrary, the native manifold $\mathcal{F}_{23}$ is still identifiable as long as the magnitude $\|\Delta p\|$ of the perturbation is not too large. Let

$$\begin{aligned}\delta \quad &= \quad \min_{\mathit{dist}\big(\,p,\,\mathcal{F}_{\ell_1\cdots\ell_k}\,\big)>0,\ \ell_1+\cdots+\ell_k=5} \mathit{dist}\,(\,p,\ \mathcal{F}_{\ell_1\cdots\ell_k}\,) \\ &= \quad \min\big\{\mathit{dist}\,(\,p,\ \mathcal{F}_{14}\,)\,,\ \mathit{dist}\,(\,p,\ \mathcal{F}_{5}\,)\,\big\}\end{aligned}$$

that is nonzero. Assume $\|\Delta p\| < \frac{\delta}{2}$. The empirical data point $\tilde{p} = p + \Delta p$ can land on any one of the manifolds $\mathcal{F}_{2\,3}$, $\mathcal{F}_{1\,1\,3}$, $\mathcal{F}_{1\,2\,2}$, $\mathcal{F}_{1\,1\,1\,2}$ and $\mathcal{F}_{1\,1\,1\,1\,1}$ but not $\mathcal{F}_{1\,4}$ or $\mathcal{F}_5$. The native manifold $\mathcal{F}_{2\,3}$ of the underlying problem p is still the manifold of the highest singularity among all the manifolds within a distance $\frac{\delta}{2}$ of the data. Namely

$$\mathit{codim}\big(\mathcal{F}_{2\,3}\big) \quad = \max_{\mathit{dist}\big(p+\Delta p, \mathcal{F}_{\ell_1\cdots\ell_k}\big)<\frac{\delta}{2},\ \ell_1+\cdots+\ell_k=5} \mathit{codim}(\mathcal{F}_{\ell_1\cdots\ell_k})$$

The problem of identifying the native manifold $\mathcal{F}_{2\,3}$ by maximizing the singularity within $\frac{\delta}{2}$ of the data point $p + \Delta p$ is clearly a well posed problem.

> *The native manifold of the underlying hypersensitive problem is identifiable from its empirical data as the manifold of the highest singularity among all manifolds with a distance below a proper error tolerance from the data point.*

Assuming the data point $\tilde{p} = p + \Delta p$ is reasonably close to p and inside the tubular neighborhood of the native manifold $\mathcal{F}_{2\,3}$, the Tubular Neighborhood Theorem ensures that there is a unique projection from $\tilde{p}$ to a point $\hat{p} = \mathbf{f}(\hat{\mathbf{z}})$ on the manifold $\mathcal{F}_{2\,3}$ and $\hat{\mathbf{z}}$ is an accurate numerical solution to the underlying problem and the solution error is in the same order of the data error.

The seven geometric observations can be rigorously proved for polynomial root-finding problem, polynomial greatest common divisor problem in both univariate and multivariate cases, multivariate polynomial factorization problem, etc. They are likely true for many other hypersensitive problems such as the problem of computing the Jordan Canonical Form of matrices.

We can reiterate those observations on the rank-revealing problem. Let A be a matrix of size, say, 8×5. The rank-revealing problem of A consists computing the rank and subspaces such as ranges and kernels. The matrix A is on the rank manifold $\mathcal{C}_r^{8\times 5}$ where r is the rank of A. The singularity $\mathit{codim}(\mathcal{C}_r^{8\times 5}) = (8-r)(5-r)$. The rank-revealing problem of A is hypersensitive if and only if the singularity is nonzero (i.e. $\mathit{rank}(A) < 5$). When A is of full-rank, the rank-revealing problem is ill-conditioned only if A is near rank-deficient matrices so that the condition number of A is large. When A is rank-deficient with $r < 5$, tiny arbitrary perturbations on A alter the rank and subspaces substantially. If the perturbation ΔA is constrained so that $A + \Delta A$ stays on the manifold $\mathcal{C}_r^{8\times 5}$, however, the rank and subspaces are not hypersensitive at all by Wedin's Theorem [12, Theorem 4.1, p. 260]. It is well-known in linear algebra that a tiny perturbation can only increase the rank and never reduce it, or equivalently can only reduces the singularity and never increase it. The rank manifolds in the vector space $\mathbb{C}^{8\times 5}$ form a simple strata

$$\begin{array}{lccccccccccc} \text{singularities:} & 40 & & 28 & & 18 & & 10 & & 4 & & 0 \\ \text{manifolds} & \mathcal{C}_0^{8\times 5} & \hookrightarrow & \mathcal{C}_1^{8\times 5} & \hookrightarrow & \mathcal{C}_2^{8\times 5} & \hookrightarrow & \mathcal{C}_3^{8\times 5} & \hookrightarrow & \mathcal{C}_4^{8\times 5} & \hookrightarrow & \mathcal{C}_5^{8\times 5}. \end{array}$$

If the rank of A is, say, 3, then A is of zero distance to manifolds $\mathcal{C}_3^{8\times 5}$, $\mathcal{C}_4^{8\times 5}$ and $\mathcal{C}_5^{8\times 5}$ and the native manifold $\mathcal{C}_3^{8\times 5}$ distinguishes itself as having the highest singularity among them. Even if A is only known thorough its empirical data as $\tilde{A} = A + \Delta A$, the native manifold $\mathcal{C}_3^{8\times 5}$ still distinguishes itself as having the highest singularity among nearby rank-manifolds.

8 Regularization of hypersensitive problems

When we face the root-finding problem of a polynomial, say $\tilde{p}(x)$ in (1), what we really want are the roots of the hidden underlying polynomial $p(x)$ in (2). with a data error bound about 2×10^{-3}. However, the polynomial p is hypersensitive for the root-finding problem in conventional sense. As shown in Figure 1, it makes little sense to solve the equation $\tilde{p}(x) = 0$ since the solutions bear almost no resemblance to the solutions we really expect.

The objective in practical computation is to *find the solution of the underlying problem accurately from empirical data* with an accuracy the data deserve. In our simulated example, the goal is to find roots of p in (2) accurately using the data polynomial such as $\tilde{p}$ in (1) with an accuracy in the same order of the data accuracy without knowing p beyond the error bound $\|p - \tilde{p}\| \leq 2 \times 10^{-3}$. For this purpose, we have no choice but to abandon the meaningless problem of "solving $\tilde{p}(x) = 0$" and modify it to a numerically feasible problem. This problem modification process for an ill-posed problem is call *regularization*.

Ideally, the regularized problem should have the identical solution at the specific unperturbed data of the original problem and the solution becomes Lipschitz continuous with respect to data in a neighborhood of the intended solution.

From the geometry, the underlying polynomial p is on the manifold $\mathcal{F}_{14}$ that is identifiable as having the highest singularity among manifolds nearby, and roots are not hypersensitive with respect to the polynomials on that manifold. We can naturally regularize the root-finding problem as calculating a specific type of *regularized roots* that are exact roots of the nearest polynomial on the factorization manifold of highest singularity passing through a neighborhood of the data polynomial.

Definition 8.1 (Regularized Roots) *Let $\tilde{p}$ be a polynomial of degree n along with an error tolerance ε. The* regularized roots *of $\tilde{p}$ within ε are exact roots of a polynomial $\hat{p}$ that is on the factorization manifold $\mathcal{F}_{\ell_1 \cdots \ell_k}$ where*

$$\begin{aligned} \mathit{codim}(\mathcal{F}_{\ell_1 \cdots \ell_k}) &= \max_{\mathit{dist}\left(\tilde{p}, \mathcal{F}_{\lambda_1 \cdots \lambda_l}\right) < \varepsilon,\ \lambda_1 + \cdots + \lambda_l = n} \mathit{codim}(\mathcal{F}_{\lambda_1 \cdots \lambda_l}) \\ \|\tilde{p} - \hat{p}\| &= \min_{u \in \mathcal{F}_{\ell_1 \cdots \ell_k}} \|\tilde{p} - u\| \end{aligned}$$

Theorem 8.2 (Fundamental Theorem of Regularized Roots) *Let p be a* *polynomial of degree* n *with roots* $z_1, \ldots, z_k$ *of multiplicities* $\ell_1, \ldots, \ell_k$ *respectively. The following properties of regularized roots hold.*

(i) The notion of regularized roots generalizes exact roots: *There is a* $\mu > 0$ *such that the regularized roots of* p *within any* $\varepsilon \in (0, \mu)$ *are identical to the exact roots of* p *with same respective multiplicities.*

(ii) The regularized root problem is well-posed: *There is a neighborhood* Ω_p *of* p *in which every polynomial* $\tilde{p}$ *serving as empirical data of* p *is associated with an interval* $(\delta_{\tilde{p}}, \theta_{\tilde{p}})$ *with* $0 \leq \delta_{\tilde{p}} \leq \|p - \tilde{p}\|$ *such that the regularized roots of* $\tilde{p}$ *within* $\varepsilon \in (\delta_{\tilde{p}}, \theta_{\tilde{p}})$ *uniquely exist as* $\tilde{z}_1, \ldots, \tilde{z}_k$ *of the same multiplicities* $\ell_1, \ldots, \ell_k$ *respectively and are Lipschitz continuous with respect to* $\tilde{p}$ *in the sense that, for every* $\tilde{q} \in \Omega_p$*, the regularized roots* $\tilde{y}_1, \ldots, \tilde{y}_k$ *of* $\tilde{q}$ *with multiplicities* $\ell_1, \ldots, \ell_k$ *respectively satisfy*

$$\big\|(\tilde{y}_1, \ldots, \tilde{y}_k) - (\tilde{z}_1, \ldots, \tilde{z}_k)\big\| \quad \leq \quad \eta_p \, \|\tilde{p} - \tilde{q}\|.$$

where η_p *is a constant depends on* p *and* ε*.*

(iii) Regularized roots are backward accurate*: For every* $\tilde{p} \in \Omega_p$*, the regularized roots of* $\tilde{p}$ *within* $\varepsilon \in (\delta_{\tilde{p}}, \theta_{\tilde{p}})$ *are exact roots of a polynomial* $\hat{p}$ *with*

$$\|\tilde{p} - \hat{p}\| \quad \leq \quad \|\tilde{p} - p\| \quad < \quad \varepsilon.$$

(iv) The regularized roots of data polynomial $\tilde{p}$ approximate the exact roots of the underlying polynomial p with an accuracy $O(\|p - \tilde{p}\|)$.

Quite obviously, the well-posedness assertion (ii) in the theorem is the corollary of the Tubular Neighborhood Theorem as established in Theorem 4.1.

This theorem is a special case of the Numerical Factorization Theorem in [14] with a rigorous and detailed proof referencing the Tubular Neighborhood Theorem in an unpublished technical report by this author. Theorem 4.1 completes the work of numerical factorization in [14] and finishes the proof of Fundamental Theorem of Regularized Roots.

This regularization follows “three-strikes principles” that have been applied effectively for striking out the hypersensitivity of singular algebraic problems:

Backward stability: The regularized numerical solution at given data is the exact solution at a specific data point within an error tolerance.

Maximum singularity: The specific data point is on a particular manifold of the highest singularity among all manifolds within the error tolerance of the given data.

Minimum distance: The specific data point is the nearest point on the particular manifold to the given data.

A hypersensitive problem such as the roots of the polynomial p in (2) is on a complex analytic manifold such as $\mathcal{F}_{41}$. Solving such problems from empirical data is possible only if the data are reasonably sound so that the data point is still in the tubular neighborhood.

9 Numerical solution of the regularized problem

In a nutshell, the geometric regularization can be summarized in one sentence:

> The regularized solution of a hypersensitive problem is the exact solution of a nearby problem that is the point on the manifold of maximum singularity with the minimum distance from the data.

A natural approach for finding the regularized solution is a two staged process:

Stage I. Within the error tolerance of the data, find the nearby complex analytic manifold of the highest singularity.

Stage II. Solve the least squares problem that minimizes the distance from the data to the manifold.

We continue using the polynomial root-finding problem for elaboration of the algorithmic engineering. Given the data polynomial $\tilde{p}$ in (1), the data error is at the third digits of the coefficients so the error tolerance can be set at 10^{-3} (more to that later).

The underlying polynomial p in (2) is on the manifold $\mathcal{F}_{41}$ if and only if the polynomial pair (p, p') belongs to the complex analytic manifold $\mathscr{P}^3_{5,4}$ since

$$\mathscr{gcd}(p, p') \;=\; (x-1)^3.$$

The singularity of $\mathscr{P}^3_{5,4}$ is identical to the degree 3 of the greatest common divisor. Write $p(x) = (x-1)^3 \cdot \big((x+1)(x+\frac{2}{3})\big)$ and

$$3 \;=\; 5 - \mathscr{deg}\big((x-1)(x+\tfrac{2}{3})\big) \;=\; n-k$$

where $n = \mathscr{deg}(p)$ and k is the number of roots. It is easy to see that finding the maximal singularity factorization polynomial near $\tilde{p}$ is equivalent to finding the maximal singularity GCD manifold near the polynomial pair $(\tilde{p}, \tilde{p}')$.

The polynomial GCD problem is also hypersensitive and can be regularized following the same three-strikes principles [18]. Using the numerical GCD finder `PolynomialGCD` in the package NACLAB[2], we can calculate the *regularized* greatest common divisor within the error tolerance 10^{-3} by the following calls:

```
>> [u,v,w,res,gcond]=PolynomialGCD(p,q,1e-3)
u =
0.89678776072183 + 2.68248022385017*x + 2.6777673064563*x^2 + 0.891138887558938*x^3

v =
0.743447961578301 + 1.86510244285362*x + 1.12216238113407*x^2
```

[2] http://homepages.neiu.edu/∼naclab

```
w =
4.08901144885262 + 5.61078724773379*x

res =
     3.122507798584681e-05

gcond =
     5.630826706824984e+02
```

The results can be interpreted as follows: The polynomial pair $(\tilde{p}, \tilde{p}')$ has a regularized greatest common divisor $\hat{u}$ and co-factors $\hat{v}$ and $\hat{w}$ such that

$$\tilde{p} \;\approx\; \hat{p} \;=\; \hat{u}\,\hat{v}, \qquad \tilde{q} \;\approx\; \hat{q} \;=\; \hat{u}\,\hat{w}.$$

The regularized co-factor $\hat{v} \approx 0.7435 + 1.8651\,x - 1.1222\,x^2$ of $\tilde{p}$ is of degree 2 with roots approximately 0.9987 and 0.6634. Furthermore, rounding up $\frac{\hat{w}(x)}{\hat{v}'(x)}$ to the nearest integer at x values -0.9987 and -0.6634 yields 4 and 1 respectively. As a result, the most singular manifold near $\tilde{p}$ within the error tolerance 10^{-3} is identified as $\mathcal{F}_{41}$ with roots approximately -0.9987 and -0.6634 of multiplicity 4 and 1 respectively, completing Stage I of the process.

Every polynomial in $\mathcal{F}_{41}$ is an image of the holomorphic mapping

$$\begin{array}{cccccc} \mathbf{f} & : & \mathbb{C}^3 & \longrightarrow & \mathbb{P}_5 \\ & & (a, z_1, z_2) & \longmapsto & a\,(x-z_1)^4(x-z_2) \end{array}$$

Thus Stage II of the computation can be carried out by solving for the least squares solution of the system

$$\mathbf{f}(a, z_1, z_2) \;=\; \tilde{p}$$

by the Gauss-Newton iteration with the initial iterate

$$(a_0, z_1^{(0)}, z_2^{(0)}) \;=\; (1, -0.9987, -0.6634).$$

The complete algorithm for computing the regularized roots is implemented in our Matlab package NAClab as the module `PolynomialFactor`. The numerical result can be obtained with a simple call:

```
>> [P,res,fcond] = PolynomialFactor(p,1e-3,'row')

P =
(1.00006771643903) * (x+0.999940850183)^4 * (x+0.666823116178)

res =
     3.624323682583038e-04

fcond =
   7.938647029187908
```

Namely, the numerical solution of the regularized factorization is

$$1.00006771643903\,(x + 0.999940850183)^4\,(x + 0.666823116178)$$

with a residual 3.62×10^{-4} and a regularized condition number 7.94 measuring the sensitivity.

The computed regularized roots are -0.99994 and -0.66682, approximating the exact roots -1 and $-\frac{2}{3}$ with the precise multiplicities 4 and 1 respectively. Considering the data error bound 10^{-3}, this is the accuracy the data deserve. The module `PolynomialFactor` also reports the residual as 3.62×10^{-4} that is the distance from the data $\tilde{p}$ to the manifold $\mathcal{F}_{4\,1}$ serving as the backward accuracy of the computing result.

Most importantly, the sensitivity measure 7.94 indicates that computing the regularized roots is highly stable even though the conventional root-finding problem is highly sensitive. As a result, we achieved a complete geometrical regularization of the ill-posed problem by applying the Tubular Neighborhood Theorem.

Acknowledgment. The author is indebted to his former colleague Marian Gidea for a casual chat in the department copy room several years ago. Knowing nothing about differential geometry beyong roughly what a manifold was, I asked: If a point is near a smooth manifold, isn't it intuitive that the projection of the point to the manifold is unique? His response: It is just the Tubular Neighborhood Theorem. The theorem is, of course, in his book [2]. That chat precipitated several papers including this one.

References

[1] P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, Princeton and Oxford, 2008.

[2] K. Burns and M. Gidea. *Differential Geometry and Topology: with a view to dynamical systems*. CRC Press, 2005.

[3] R. M. Corless, A. Galligo, I. Kotsireas, and S. Watt. A geometric-numeric algorithm for factoring multivariate polynomials. Proc. ISSAC'02, ACM Press, pages 37-45, 2002.

[4] J. W. Demmel and A. Edelman. The dimension of matrices (matrix pencils) with given Jordan (Kronecker) canonical forms. *Linear Algebra and its Applications*, 230:61–87, 1995.

[5] J. A. Dieudonné. *A History of Algebraic and Differential Topology, 1900-1960.* Birkhäuser, Boston, 1989.

[6] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonality constraints. *SIAM J. Matrix Anal. Appl.*, 20:303–353, 1998.

[7] A. Edelman, E. Elmroth, and B. Kågström. A geometric approach to perturbation theory of matrices and and matrix pencils. Part I: Versal deformations. *SIAM J. Matrix Anal. Appl.*, 18:653–692, 1997.

[8] A. Edelman, E. Elmroth, and B. Kågström. A geometric approach to perturbation theory of matrices and and matrix pencils. Part II: a stratification-enhanced staircase algorithm. *SIAM J. Matrix Anal. Appl.*, 20:667–699, 1999.

[9] W. Kahan. Conserving confluence curbs ill-condition. Technical Report 6, Computer Science, University of California, Berkeley, 1972.

[10] R. A. Lippert and A. Edelman. The computation and sensitivity of double eigenvalues. In *Advances in computational mathematics, Lecture Notes in Pure and Appl. Math. 202*, pages 353–393, New York, 1999. Dekker.

[11] A. J. Sommese, J. Verschelde, and C. W. Wampler. Introduction to numerical algebraic geometry. In A. Dickenstein and I. Z. Emiris, editors, *Solving Polynomial Equations*, pages 301–337, Springer-Verlag Berlin Heidelberg, 2005.

[12] G. W. Stewart and J. guang Sun. *Matrix Perturbation Theory.* Academic Press, Inc, Boston, San Diego, New York, London, Sydney, Tokyo, Toronto, 1990.

[13] N. Trefethen. Maxims about numerical mathematics, science, computers, and life on earth. *SIAM News*, 31(1):4, 1998.

[14] W. Wu and Z. Zeng. The numerical factorization of polynomials. *J. Foundation of Computational Mathematics*, 2016. appear, DOI 10.1007/s10208-015-9289-1.

[15] Z. Zeng. Algorithm 835: MultRoot – A Matlab package for computing polynomial roots and multiplicities. *ACM Trans. Math. Software*, 30:218–235, 2004.

[16] Z. Zeng. Computing multiple roots of inexact polynomials. *Math. Comp.*, 74:869–903, 2005.

[17] Z. Zeng. The approximate irreducible factorization of a univariate polynomial. Revisited. Proceedings of ISSAC '09, ACM Press, pp. 367–374, 2009.

[18] Z. Zeng. The numerical greatest common divisor of univariate polynomials. In J. R. L. Gurvits, P. Pébay and D. Thompson, editors, *Contemporary Mathematics Vol. 556, American mathematical Society,* Randomization, Relaxation and Complexity in Polynomial Equation Solving, pages 187–217, Providence, RI, 2011.