

NOT FOR QUOTATION
WITHOUT THE PERMISSION
OF THE AUTHOR

**SINGULARITY THEORY FOR NONLINEAR
OPTIMIZATION PROBLEMS**

J. Casti

November 1985
WP-85-79

Working Papers are interim reports on work of the International Institute for Applied Systems Analysis and have received only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute or of its National Member Organizations.

INTERNATIONAL INSTITUTE FOR APPLIED SYSTEMS ANALYSIS
2361 Laxenburg, Austria

Abstract

Techniques from the theory of singularities of smooth mappings are employed to study the reduction of nonlinear optimization problems to simpler forms. It is shown how singularity theory ideas can be used to : 1) reduce decision space dimensionality; (2) transform the constraint space to simpler form for primal algorithms; (3) provide sensitivity analysis.

SINGULARITY THEORY FOR NONLINEAR OPTIMIZATION PROBLEMS

J. Casti

I. Background

Consider a smooth (C^∞) function $f: R^n \rightarrow R^m$ and assume that f has a critical point at the origin, i.e. $df(0) = 0$. The theory of singularities as developed by Thom, Mather, Arnol'd and others [1-3] addresses the following basic questions:

A. What is the local character of f in a neighborhood of the critical point? Basically, this question amounts to asking "at what point is it safe to truncate the Taylor series for f ?" This is the *determinacy* problem.

B. What are the "essential" perturbations of f ? That is, what perturbations of f can occur which change the qualitative nature of f and which cannot be transformed away by a change of coordinates? This is the *unfolding* problem.

C. Can we classify the types of singularities which f can have up to diffeomorphism? This is the *classification* problem.

Elementary catastrophe theory largely solves these three problems (when $m=1$); its generalization to singularity theory solves the first two, and gives relatively complete information on the third for m, n small. Here we outline a program for the utilization of these results in an applied setting to deal with certain types of nonlinear optimization problems. In the following section we give a brief summary of the main results of singularity theory for problems A-C for *functions* ($m=1$) and then proceed to a discussion of how these results may be employed for nonlinear optimization.

II. Determinacy, Unfoldings, and Classifications

Equivalence of Germs

In its local version, elementary catastrophe theory deals with functions $f: U \rightarrow R$ where U is a neighborhood of O in R^n . The cleanest way to handle such functions is to pass to *germs*, a germ being a class of functions which agree on suitable neighborhoods of O . All operations on germs are defined by performing similar operations on representatives of their classes. In the sequel, we shall usually make no distinction between a germ and a representative function.

We let E_n be the set of all smooth germs $R^n \rightarrow R$, and let E_{nm} be the set of all smooth germs $R^n \rightarrow R^m$. Of course $E_{n,1} = E_n$. These sets are vector spaces over R , of infinite dimension. We abbreviate $(x_1, \dots, x_n) \in R^n$ to x . If $f \in E_{nm}$ then

$$f(x) = (f_1(x), \dots, f_m(x))$$

and the f_i are the *components* of f .

A diffeomorphism germ $\varphi: R^n \rightarrow R^n$ satisfies $\varphi(0) = 0$, and has an inverse φ' such that $\varphi(\varphi'(x)) = x = \varphi'(\varphi(x))$ for x near 0. It represents a smooth, invertible local coordinate change. By the Inverse Function Theorem, φ is a diffeomorphism germ if and only if it has a nonzero Jacobian, that is,

$$\det[\partial\varphi_i / \partial x_j(0)] \neq 0.$$

Two germs $f, g: R^n \rightarrow R$ are *right equivalent* if there is a diffeomorphism germ φ and a constant $\gamma \in R$ such that

$$g(x) = f(\varphi(x)) + \gamma.$$

This is the natural equivalence for studying topological properties of the gradient ∇f (Poston and Stewart [4]). If f , rather than ∇f , is important, the term γ is omitted.

A *type* of germ is a right equivalence class and the classification of germs up to right equivalence amounts to a classification of types. Each type forms a subset of E_n , and the central object of study is the way these types fit together.

A precise description is complicated by the fact that most types have infinite dimension; but there is a measure of the complexity of a type, the *codimension*, which is generally finite. Heuristically, it is the difference between the dimension of the type and that of E_n (even though both are infinite). A precise definition is given below.

The largest types have codimension 0 and form open sets in E_n . Their boundaries contain types of codimension 1; the boundaries of these in turn contain types of codimension 2, and so on, with higher codimensions revealing progressively more complex types. Types of infinite codimension exist, but form a very small set in a reasonable sense.

Codimension and the Jacobian ideal

Let E_n be the set of germs $R^n \rightarrow R$, and let F be the set of formal power series in x_1, \dots, x_n . There is a map $j: E \rightarrow F$ defined by

$$jf = f(0) + \sum \frac{\partial f}{\partial x_i}(0)x_i + \frac{1}{2} \sum \frac{\partial^2}{\partial x_i \partial x_j}(0)x_i x_j + \dots$$

where the right-hand side is the Taylor series, or *jet*, of f . Note that it exists as a *formal* power series for all smooth f : convergence is not required in what follows. The map j is onto, linear over R , and preserves products (i.e., $j(f \cdot g) = j(fg) = (jf \cdot jg)$).

Let m_n be the set of $f \in E_n$ such that $f(0) = 0$. This is an *ideal* of E_n (meaning that if $f \in m_n$ and $g \in E_n$ then $fg \in m_n$, which we write briefly as $m_n E_n \subset m_n$). Its k th power m_n^k consists of all $f \in E_n$ such that $0 = f(0) = df(0) = d^2 f(0) = \dots = d^{k-1} f(0)$. In particular, f is a *singularity* if and only if $f \in m_n$. The ideals m_n^k form a decreasing sequence.

$$E_n \supset m_n \supset m_n^2 \supset m_n^3 \supset \dots$$

There is a similar chain in F_n . Let $M_n = j(m_n)$: this is the set of formal power series with zero constant term. Then $M_n^k = j(m_n^k)$ is the set of formal power series without terms of degree $\leq k-1$. The intersection of all M_n^k is 0; the intersection of all m_n^k is the set m_n^∞ of flat germs, having zero Taylor series.

The *Jacobian ideal* of a singularity f is the set of all germs expressible in the form

$$g_1 \frac{\partial f}{\partial x_1} + \dots + g_n \frac{\partial f}{\partial x_n}$$

for arbitrary germs g_i . We denote it by $\Delta(f)$, or merely Δ when f is understood. Its image $j\Delta(f) \subseteq F_n$ has an analogous definition, where the partial derivatives are defined formally. Since f is a singularity, $\Delta(f) \subseteq \mathfrak{m}_n$. The *codimension* of f is defined to be

$$\text{cod}(f) = \dim_R \mathfrak{m}_n / \Delta(f).$$

Similarly, at the formal power series level, we define

$$\text{cod}(jf) = \dim_R M_n / j\Delta(f).$$

The codimension of an orbit is the same as that of its tangent space T . This is the same as the dimension of the quotient vector space E/T . In E_n , the analog of this tangent space is the Jacobian ideal, so the codimension should be $\dim E_n / \Delta(f)$. This measures the number of independent directions in E_n "missing" from $\Delta(f)$, or equivalently missing from the orbit of f .

The computation of $\text{cod}(f)$ is effected by means of the following result: if either $\text{cod}(f)$ or $\text{cod}(jf)$ is finite then so is the other, and they are equal. Thus, the computation may be carried out on the formal power series level where it is a combinatorial calculation. For examples in classical notation, see Poston and Stewart [4].

Determinacy

Let $f \in E_n$, and define the k -jet $j^k(f)$ to be the Taylor series of f up to and including terms of order k . For example,

$$j^6(\sin(x)) = x - x^3/3! + x^5/5!.$$

We say that f is k -*determinate* (or k -determined) if for any $g \in E_n$ such that $j^k g = j^k f$, it follows that g is right equivalent to f .

A germ is 1-determined if its linear part is nonzero, that is, its derivative does not vanish. So the non-1-determined germs are the singularities. If f is a singularity and $f(0) = 0$ (as we can assume) then the second derivative gives the 2-jet of f in the form

$$j^2 f(x_1, \dots, x_n) = \sum_{i,j} H_{ij} x^i x^j$$

where the Hessian matrix

$$H = (H_{ij})(0)$$

is symmetric. It can be shown that f is 2-determined if and only if $\det(H) \neq 0$; in this case f is right equivalent to

$$\pm x_1^2 \pm \dots \pm x_n^2. \quad (*)$$

This is a reformulation in determinacy terms of the *Morse Lemma* (Milnor [5]). A germ equivalent to (*) is said to be *Morse*. Morse germs are precisely those of codimension 0. The number l of negative signs in (*) is the *index* of f , and f is an l -*saddle*. Morse theory (Milnor [5]) describes the *global* properties of a function

$f: X \rightarrow R$ where X is a smooth manifold, and f has only Morse singularities. (See Casti [9] for more details).

There exist rules for computing the determinacy of a given germ: an easy necessary condition, an easy (different) sufficient condition, and a harder necessary-and-sufficient condition.

Let Δ be the Jacobian ideal of f . Then:

- (i) If $m_n^k \subset m_n \Delta$ then f is k -determined.
- (ii) If f is k -determined then $m_n^{k+1} \subset m_n \Delta$.
- (iii) f is k -determined if and only if $m_n^{k+1} \subset m_n \Delta(f+g)$ for all $g \in m_n^{k+1}$.

There is a slightly stronger form of (i), namely

- (i') If $m_n^{k+1} \subset m_n^2 \Delta$ then f is k -determined.

Numerous examples in Poston and Stewart [4] and Gibson [3] show how to compute the determinacy of a given f . For example, suppose f is in Morse form (*). Then $\Delta = \langle +2x_1, \dots, +2x_n \rangle = m_n$ and $m_n^2 = m_n \Delta$. By (i), f is 2-determined as asserted above.

A germ is *finitely* determined if it is k -determined for some finite k . The following are equivalent:

- (iv) f has finite codimension
- (v) f is finitely determined
- (vi) $m_n^t \subset \Delta$ for some t .

The solution to the Determinacy Problem is thus that it is safe (up to right equivalence) to truncate a k -determined germ at degree k of its Taylor series. For a germ such as $x^2y \in E_2$, which is not finitely determined, it is not safe to truncate higher order perturbing terms (and indeed $x^2y + y^t$ has a type that depends on t). Germs that are not finitely determined either arise in a context where some symmetry is acting (and should be analyzed by methods similar to those above but which take symmetry into account - which can be done) or must be viewed with suspicion. By (iv), we may summarize: "nice" germs have finite codimension.

Suppose that f is not 2-determinate, so that $\det(H)=0$. Let the rank of the matrix H be r and call $n-r$ its *corank*. A useful result, called the *Splitting Lemma*, says that f is right equivalent to a germ of the form

$$g(x_1, \dots, x_{n-r}) \pm x_{n-r+1}^2 \pm \dots \pm x_n^2.$$

For many purposes, the quadratic terms may be ignored. So the Splitting Lemma reduces the effective number of variables to $n-r$. A simple proof for finite dimensions is in Poston and Stewart [4].

The determinacy calculations, and the application of the Splitting Lemma, may be carried out equally well on $j^k f$ in F_n , provided the codimension of f is finite. The formal power series setting is better for computations.

Unfoldings

An unfolding of a singularity is a "parametrized family of perturbations." The notion is useful mainly because, for finite codimension singularities, there exists a "universal unfolding" which in a sense captures all possible unfoldings.

More rigorously, let $f \in E_n$. Then an l -parameter unfolding of f is a germ $F \in F_{n+l}$, that is, a real-valued germ of a function $F(x_1, \dots, x_n, \varepsilon_1, \dots, \varepsilon_l) = F(x, \varepsilon)$, such that $F(x, 0) = f(x)$.

An unfolding F is *induced* from f if

$$F(x, \delta) = F(\rho_\delta(x), \psi(\delta)) + \gamma(\delta)$$

where

$$\delta = (\delta_1, \dots, \delta_m) \in R^m$$

$$\rho_\delta: R^n \rightarrow R^n$$

$$\psi: R^m \rightarrow R^l$$

$$\gamma: R^l \rightarrow R.$$

Two unfoldings are *equivalent* if each can be induced from the other. An l -parameter unfolding is *versal* if all other unfoldings can be induced from it; *universal* if in addition, l is as small as possible.

Suppose that f has finite codimension c . Let u_1, \dots, u_c be a basis for $\mathfrak{m}_n / \Delta(f)$. Then it is a theorem that a *universal unfolding* is given by the germ

$$F(x, \varepsilon) = f(x) + \varepsilon_1 u_1(x) + \dots + \varepsilon_c u_c(x), \varepsilon_i \in R \quad (**)$$

While different choices of the u_i can be made, a universal unfolding is unique up to equivalence. The existence of universal unfoldings in finite codimension, and the method for computing them, is probably the most significant and useful result in elementary catastrophe theory. (Note that (**) is linear in the unfolding variables ε . This is a theorem, and is *not* built into the definition of an unfolding.)

For example, if $f(x, y) = x^3 + y^4$, then a basis for $\mathfrak{m}_2 / \Delta(f)$ is $\{x, y, xy, y^2, xy^2\}$. So a universal unfolding is given by

$$F(x, y, \varepsilon) = x^3 + y^4 + \varepsilon_1 x + \varepsilon_2 y + \varepsilon_3 xy + \varepsilon_4 y^2 + \varepsilon_5 xy^2.$$

The codimension of a germ f has several interpretations:

- (i) The codimension of the Jacobian ideal in \mathfrak{m}_n ,
- (ii) The number of independent directions "missing" from the orbit of f ,
- (iii) The number of parameters in any universal unfolding of f .

In addition, if the codimension of f is c , it can be shown that any small perturbation of f has at most $c+1$ critical points.

Classification

We sketch how these ideas may be used to classify germs of codimension at most 4.

Let $f \in E_n$. If f is not a singularity then $f(x)$ is right equivalent to x_1 . If f is a singularity, and its Hessian has nonzero determinant, then f is right equivalent to $\pm x_1^2 \pm \dots \pm x_n^2$. Otherwise, $\det(H) = 0$. Let $k = n - r$ be the corank of H , and split f as

$$f(x) = g(x_1, \dots, x_k) \pm x_{k+1}^2 \pm \dots \pm x_n^2.$$

It can be proved that the classification of possibilities for f depends only on the similar classification for g .

The Taylor series of g begins with cubic or higher terms. First suppose that $k=1$, and let the first nonzero jet of g be $a_t x^t$. This is t -determined, and scales to $\pm x^t$ (t even), x^t (t odd). The codimension is $t-2$, so $t = 3, 4, 5$ or 6 .

Next, let $k=2$, and let

$$j^3 g(x, y) = ax^3 + bx^2y + cxy^2 + dy^3.$$

By a linear change of variable, this cubic may be brought to the form $x^3 + xy^2$ (one real root), $x^3 - xy^2$ (three distinct real roots), x^2y (three real roots, one repeated), x^3 (three real roots, all repeated), or 0 .

The forms $x^3 \pm xy^2$ are 3-determined, and of codimension 3.

The form x^2y is not 3-determined, so we consider higher terms. A series of changes of variable bring any higher order expansion to the form $x^2y + y^t$, which is t -determined and of codimension t . Only $t=4$ is relevant to our problem here.

No higher term added to x^3 produces a codimension 4 result; and no higher term added to 0 does.

Finally, let $k \geq 3$. Then the codimension can be proved to be at least 7, so this case does not arise.

Thus, we have classified the germs of codimension ≤ 4 into the canonical forms

$$\begin{aligned} & x_1 \\ & \pm x_1^2 \pm \cdots \pm x_n^2 \\ & x_1^3 + (M) \\ & x_1^4 + (M) \\ & x_1^5 + (M) \\ & x_1^6 + (M) \\ & x_1^3 - x_1 x_2^2 + (N) \\ & x_1^3 + x_1 x_2^2 + (N) \\ & x_1^3 + x_2^4 + (N) \end{aligned}$$

where

$$(M) = \pm x_2^2 + \cdots \pm x_n^2, (N) = \pm x_3^2 \pm \cdots \pm x_n^2.$$

The celebrated elementary catastrophes of Thom are the universal unfoldings of the singularities on this list, or its extension to higher codimensions. The universal unfolding arises when we try to classify not germs, but l -parameter families of germs. For $l \leq 4$, "almost all" such are given by universal unfoldings of germs of codimension ≤ 4 .

Table 1 summarizes the list of germs and their unfoldings up to codimension 5, together with their customary name and symbol in the systematic notation of Arnol'd [1]. The terms (M) and (N) are omitted for clarity, x and y replace x_1 and x_2 ; and unfolding parameters are listed as (a,b,c,d,e) rather than $(\varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, \varepsilon_5)$.

TABLE 1: The elementary catastrophes of codimension ≤ 5 . When the \pm sign occurs, germs with sign (+) are called *standard*, (-) are called *dual*.

symbol	name	germ	universal unfolding	co-rank	codimension
A_2	fold	x^3	x^3+ax	1	1
$\pm A_3$	cuspid	$\pm x^4$	$\pm x^4+ax^2+bx$	1	2
A_4	swallowtail	x^5	$x^5+ax^3+bx^2+cx$	1	3
$\pm A_5$	butterfly	$\pm x^6$	$\pm x^6+ax^4+bx^3+cx^2+dx$	1	4
A_6	wigwag	x^7	$x^7+ax^5+bx^4+cx^3+dx^2+ex$	1	5
D_4^-	elliptic umbilic	x^3-xy^2	$x^3-xy^2+ax^2+bx+cy$	2	3
D_4^+	hyperbolic umbilic	x^3+xy^2	$x^3+xy^2+ax^2+bx+cy$	2	3
$\pm D_5$	parabolic umbilic	$\pm(x^2y+y^4)$	$\pm(x^2y+y^4)+ax^2+by^2+cx+dy$	2	4
D_6^-	second elliptic umbilic	x^5-xy^2	$x^5-xy^2+ay^3+bx^2+cy^2+dx+ey$	2	5
D_6^+	second hyperbolic umbilic	x^5+xy^2	$x^5+xy^2+ay^3+bx^2+cy^2+dx+ey$	2	5
$\pm E_6$	symbolic umbilic	$\pm(x^3+y^4)$	$\pm(x^3+y^4)+axy^2+by^2+cxy+dx+ey$	2	5

The above sketch shows how the classification problem reduces to the determinacy and unfolding problems (and is relatively easy once these are solved). In applications, the main influence of the classification is an organizing one: the determinacy and unfolding theorems play a more direct role.

III. Singularity Theory and Nonlinear Programming

We consider the problem

$$\max f(x) \tag{1}$$

over all $x \in \mathbb{R}^n$ such that

$$g(x) \leq 0, \tag{2}$$

where $f, g \in \mathcal{M}_n$. There are at least three different aspects of this standard nonlinear optimization problem which singularity theory can shed some light upon: (1) reduction of dimensionality in the decision space for dual, penalty, and barrier type algorithms [7]; (2) transformation of the constraint space into simpler form for primal type algorithms [7] and (3) sensitivity analysis. Let us examine each of these areas in turn.

Dimensionality Reduction and the Splitting Lemma

If the optimization problem (1)-(2) is to be approached using one of the dual penalty, or barrier algorithms [7], the Splitting Lemma can be used to reduce the dimension of the decision vector in the surrogate objective function. For example, consider the augmented Lagrangian method, for which the surrogate objective function is

$$G(x, \alpha) = \alpha'g + f + \rho/2 \|g\|^2,$$

where α is a vector of multipliers and ρ is some positive constant. The parameters α are updated according to, say, the augmented Lagrangian scheme of Hestenes.

Assume that the critical point of G is located at $x = x^*, \alpha = \alpha^*$, and that the corank of $G(x, \alpha) = r$. Then the Splitting Lemma insures that there exist coordinate transformations $x \rightarrow \hat{x}, \alpha \rightarrow \hat{\alpha}$ such that $G \rightarrow \hat{G}$, where

$$\hat{G}(\hat{x}, \hat{\alpha}) = G_1(\hat{x}_1, \dots, \hat{x}_r, \hat{\alpha}_1, \dots, \hat{\alpha}_c) + M(x_{r+1}, \dots, x_n),$$

where $c = \text{codim } G$ while $G_1(\cdot)$ is a function $O(\|x\|^3)$, which is linear in $\hat{\alpha}_1, \dots, \hat{\alpha}_c$. The function $M(\cdot)$ is a pure quadratic. The important point here is that usually $r \ll n$, which implies that most of the computational work is involved in minimizing the quadratic M , which can be done very efficiently by any of a number of quasi-Newton schemes. The essentially nonlinear part of the problem involves the minimization of G , which, however, involves only r variables. Often $r = 1$ or 2 , even if n is very large, say, hundreds, so the computational savings can be significant.

The potential drawback to the above scheme is that in order to compute r , the corank of G , we need to know the Hessian

$$H = \left[\frac{\partial^2 G}{\partial x^2} \right]$$

at the critical point (x^*, α^*) . Since it is precisely x^* which we seek, it appears at first glance that the situation is not too promising. However, this problem can be circumvented in at least two different ways:

(i) Often it can be seen that the Hessian will be of constant rank in some neighborhood D of x^* , even if we don't know x^* exactly. This situation comes about since we usually have at least some idea of the region D containing x^* . Thus, if we have an estimate of D and know that $\text{rank } H(x, \alpha) = \text{constant}$ for all $x \in D$, then we can use this information in a successive approximation scheme generating a sequence $x_n \rightarrow x^*$. The idea is to apply the Splitting Lemma to each approximate problem at the point x_n .

(ii) if there is no information about the rank H , then we can appeal to the inequality

$$r(r+1)/2 \leq \text{codim } G,$$

which always holds. We can take a pessimistic estimate of r which, at worst, means only that we include a few more variables in our nonlinear optimization of $G_1(\cdot)$ than might have been needed. If $\text{codim } G \leq 2$, then we can see from the inequality that $r=1$ and there is only a single essential, nonlinear variable, *regardless* of where \mathbf{x}^* is located. Otherwise there may be several nonlinear variables, but the number will still be severely limited by the above inequality.

An essential ingredient in making the above scheme work in practice is the ease of determining the coordinate transformations $\mathbf{x} \rightarrow \hat{\mathbf{x}}, \alpha \rightarrow \hat{\alpha}$. As noted in Section II, the theory guarantees such transformations exist and, moreover, that they are themselves diffeomorphisms. Thus, the coordinate changes

$$\begin{aligned}\hat{x}_i &= \hat{x}_i(x_1, x_2, \dots, x_n), \\ \hat{\alpha}_j &= \hat{\alpha}_j(\alpha_1, \alpha_2, \dots, \alpha_m)\end{aligned}$$

have convergent power series expansions. Consequently, since we know the original form of G and its normal form \hat{G} , in principle we can substitute the above expansions and match coefficients in order to determine the explicit form of the transformations. The operational implementation of this idea, however, may require a substantial amount of algebra, depending upon the exact nature of G .

Simplifying the Constraint Space

For nonlinear constrained optimization problems having nonlinear constraint sets, the coordinate changes discussed above can be employed to "straighten-out" the binding constraints in a neighborhood of regular points, so that primal methods for solving constrained optimization problems can be used, dealing only with *linear* side constraints. The essence of the primal methods is to start with a feasible direction along which the objective function is improving. A one-dimensional line search (interval bisection, Newton's method, etc.) is then used to solve the one-dimensional optimization problem along the improving feasible direction, constrained so that the resulting solution remains feasible [7].

A specific example of such a primal method is the *gradient projection* technique due to Rosen. This method generates an improving feasible direction by projecting the negative of the gradient vector of f onto the affine subspace determined by the intersection of the binding constraints, assuming the constraints are *linear*. A projection matrix P is formed from a suitable linear combination of the normal vectors of the constraint subspaces (i.e. the gradients of the binding constraints). The resulting one-dimensional optimization is then guaranteed to remain feasible as long as a suitable upper bound is observed on the line search [7].

In the event the constraints are nonlinear, the gradient of f is projected onto the intersection of the tangent spaces to the binding constraints, so that movement along the improving feasible direction will, in general, take the solution outside the feasible region (see Figure 1). This necessitates a correction move to bring the solutions back into the feasible regions after the one-dimensional search has been completed. Singularity theory appears to offer the possibility of materially improving the above procedure as we now indicate.




FIGURE 1 Projected gradient method of Rosen for nonlinear constraints (From Figure 10.5, pg. 398 Bazarra and Shetty, 1979)

Consider the following nonlinear programming problem:

$$\begin{aligned} \text{minimize:} & \quad f(x) \\ \text{subject to:} & \quad g_i(x) \leq 0 \quad i=1,2,\dots,m \\ & \quad x \geq 0 \end{aligned}$$

For any x such that $x \geq 0$, if $I = \{i: g_i(x) = 0\}$, then

$$X = \{x: g_i(x) = 0\} = \bigcap_{i \in I} (g_i(x) \cap R^{n-1} \text{ hyperplane})$$

will be the intersection of a finite number of manifolds in R^n and thus, with the possible exclusion of a set of points of codimension n , (corners) will inherit the manifold structure locally. Locally, then, a coordinate change could be effected in X which will cause X to take the form:

$$X \rightarrow Y = \{y: 0 = a'y + c, \quad a, c \text{ constant vectors}\}$$

as long as the gradients of the binding constraints don't vanish. A transversality argument can be used to rule out the latter possibility.

Assuming that only the constraints $g_i(x) = 0$ is binding, let

$$S_i = T_x g_i(x) \cap R^{n-1} \text{ hyperplane,}$$

where

$$T_x g_i(x) = \text{tangent space to } g_i \text{ at } x.$$

Since $\text{codim } T_x g_i(x) = 1$ and $\text{codim } \{R^{n-1} \text{ hyperplane}\} = 1$, if the intersection is transverse

$$\text{codim } T_x g_i(x) + \text{codim } \{\mathcal{R}^{n-1} \text{ hyperplane}\} = \text{codim } S_i = 2$$

Results from differential topology assert that the set of critical points R_i for g_i will be isolated, thus the $\dim R_i = 0$ and $\text{codim } R_i = n$. Therefore,

$$\text{codim } R_i + \text{codim } S_i = n + 2 > n.$$

So, for $\nabla g_i(x)$ to be zero at exactly the same points where $g_i(x)=0$ constitutes a non-transverse intersection and is therefore non-generic. If any such points should occur, they will be isolated and thus not form a constraint boundary.

In practice, finding X and the coordinate transformation necessary to make it look like Y usually requires some effort. However, if projection onto only one binding constraint is necessary, the calculation becomes simpler, as the following example shows:

$$\min f(x_1, x_2) = 1/2x_1^2 + 1/2x_2^2 - x_1 - x_2$$

(the geometry in x space is shown in Figure 2)

subject to

$$\begin{aligned} x_1^2 + x_2^2 - 1 &\leq 0 \\ -x_1 &\leq 0 \\ -x_2 &\leq 0 \end{aligned}$$

$\nabla f(x) = (x_1 - 1, x_2 - 1)$	at (1,0):	$\nabla f(1,0) = (0, -1)$
$\nabla g_1(x) = (2x_1, 2x_2)$		$\nabla g_1(1,0) = (2,0)$ - binding
$\nabla g_2(x) = (-1, 0)$		$\nabla g_2(1,0) = (-1,0)$
$\nabla g_3(x) = (0, -1)$		$\nabla g_3(1,0) = (0, -1)$ - binding

As can be seen, we want to project onto $g_1(x)$. To straighten out g_1 , let $y_1 = x_1^2, y_2 = x_2^2$. In the new coordinates, ∇f_{new} will be :

$$\nabla f_{\text{new}}(y) = (y_1^{-1/2} - 1, y_2^{-1/2} - 1), \quad \nabla f_{\text{new}}(1,0) = (0, -1)$$

(Note: This is not the gradient of the transformed objective function but rather the transformed gradient of the old objective function.)

The new problem is :

$$\min f(y_1, y_2) = 1/2 y_1 + 1/2 y_2 - y_1^{-1/2} - y_2^{-1/2}$$

(the geometry in y space is shown in Figure 3)

subject to

$$\begin{aligned} y_1 + y_2 - 1 &\leq 0 \\ -y_1 &\leq 0 \\ -y_2 &\leq 0 \end{aligned}$$

Now the constraint is linear and we project ∇f_{new} onto g_1 by forming the projection matrix:




FIGURE 2 Configuration in x space

$$\nabla g_1 = M = (1 \ 1), \quad MM^t = (1 \ 1) \begin{bmatrix} 1 \\ 1 \end{bmatrix} = 2 \quad (MM^t)^{-1} = (1/2)$$

$$P = I - M^t(MM^t)^{-1}M = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \end{bmatrix} (1/2) (1 \ 1) = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

$$d = -P\nabla f_{\text{new}} = - \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix} \begin{bmatrix} 0 \\ -1 \end{bmatrix} = \begin{bmatrix} -\frac{1}{2} \\ \frac{1}{2} \end{bmatrix}$$

The objective function is optimized along the constraint by letting

$$y = \begin{bmatrix} 1 \\ 0 \end{bmatrix} + \begin{bmatrix} -\frac{1}{2} h \\ \frac{1}{2} h \end{bmatrix}$$

$$f(h) = 1/2(1 - 1/2 h) + 1/4 h - (1 - 1/2 h)^{1/2} - (1/2 h)^{1/2}$$

$$\frac{df}{dh} = 1/4(1 - 1/2 h)^{-1/2} - 1/4(1/2 h)^{-1/2} = 0 \Rightarrow h = 1$$

So the minimum is taken on at

$$y = \begin{bmatrix} \frac{1}{2} \\ \frac{1}{2} \end{bmatrix} \text{ or } x = \begin{bmatrix} \frac{1}{\sqrt{2}} \\ \frac{1}{\sqrt{2}} \end{bmatrix}$$

That this is the optimum can be seen by trying to form an improving feasible direction in x space. The result will be the zero vector, indicating that the optimum has been reached.

$$\nabla g_1(x) = \begin{bmatrix} \frac{2}{\sqrt{2}} & \frac{2}{\sqrt{2}} \end{bmatrix} MM^t = (\sqrt{2} \ \sqrt{2}) \begin{bmatrix} \sqrt{2} \\ \sqrt{2} \end{bmatrix} = 4 \quad (MM^t)^{-1} = 1/4$$

$$\nabla f(x) = (\sqrt{2}-1, \sqrt{2}-1) \frac{1}{\sqrt{2}} P = I - M^t(MM^t)^{-1}M = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} - \begin{bmatrix} \sqrt{2} \\ \sqrt{2} \end{bmatrix} \frac{1}{4} (\sqrt{2} \ \sqrt{2}) = \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

$$d = P\nabla f(x) = \frac{1}{\sqrt{2}} \begin{bmatrix} \frac{1}{2} & -\frac{1}{2} \\ -\frac{1}{2} & \frac{1}{2} \end{bmatrix} \begin{bmatrix} \sqrt{2}-1 \\ \sqrt{2}-1 \end{bmatrix} = 0$$




FIGURE 3 Configuration in y space

as claimed.

A summary of the algorithm is given as follows:

Initialization step: Choose a feasible point \mathbf{x}_u and find $I_1 = \{i: g_1(\mathbf{x}) = 0\}$. Let $\mathbf{u} = 1$ and go to (1).

(1) If $I_i = 0$, let $P = I$, form $\mathbf{d}_u = P \nabla f(\mathbf{x}_u)$ and go to (3). Otherwise, form the projection matrix in \mathbf{x} -space as follows. Let $M = D_g(\mathbf{x}_u)$ be the matrix of gradients of the binding constraints at \mathbf{x}_u . If $P = I - M^t(MM^t)^{-1}M = 0$, let $W = -(MM^t)^{-1}M \nabla f(\mathbf{x}_u)$. If $W \geq 0$, \mathbf{x}_u will be a Kuhn-Tucker point, otherwise, delete a row corresponding to $W_i \geq 0$ and repeat step (1). This has the effect of eliminating binding constraints from consideration which won't generate an improving feasible direction. Let $I = \{i: g_i(\mathbf{x}) = 0\}$ after a nonzero \hat{P} has been found.

(2) If $X = \bigcap_{i \in I} (g_i(\mathbf{x}) \cap \{R^{n-1} \text{ hyperplane}\})$ is already linear, use the matrix \hat{P} in the following calculations. Otherwise, find a coordinate change such that X becomes

$$Y = \{\mathbf{y}: \mathbf{a}^t \mathbf{y} + c = 0\}$$

Find $\nabla f_{\text{new}}(\mathbf{y}(\mathbf{x})) | \mathbf{y}(\mathbf{x}_u)$ and convert the problem into \mathbf{y} coordinates. Form the projection matrix $P = I - \mathbf{a}^t(\mathbf{a}\mathbf{a}^t)^{-1}\mathbf{a}$ and go to (3) after forming $\mathbf{d}_u = P \nabla f_{\text{new}}(\mathbf{y}(\mathbf{x})) | \mathbf{y}(\mathbf{x}_u)$

(3) Let \mathbf{h}_u be a solution to

$$\begin{aligned} \text{Minimize } f(\mathbf{z}_u + h\mathbf{d}_u), \text{ where } \mathbf{z}_u = \mathbf{x}_u, \text{ if in } \mathbf{x} \text{ coordinates and} \\ \mathbf{z}_u = \mathbf{y}_u, \text{ if in } \mathbf{y} \text{ coordinates} \end{aligned}$$

$$0 \leq h \leq h_{\max}$$

where h_{\max} is determined so that the problem remains feasible.

Let $\mathbf{z}_{u+1} = \mathbf{z}_u + h\mathbf{d}_u$, convert to \mathbf{x} coordinates, if necessary, and return to (1).

For more complex problems involving more than one binding constraint, the coordinate changes must be automated and checks made on the neighborhood of validity of the transformations. Application to other primal methods can also be made using the same types of arguments.

Sensitivity Analysis and Unfoldings

In Section II, we noted that a universal unfolding of a smooth function $f(\mathbf{x})$ represents the most general type of smooth perturbation to which f can be subjected and that the number of terms needed to characterize all such perturbations equals $\text{codim } f$. Furthermore, if $\mathbf{u}_1(\mathbf{x}), \dots, \mathbf{u}_c(\mathbf{x})$ represent a basis for the Jacobian ideal $\mathfrak{m}_n / \nabla(f)$, then the $\{\mathbf{u}_i\}$ also represent a basis for the space of all such perturbations. Since perturbations in the objective function and/or constraints lie at the heart of sensitivity analysis for nonlinear optimization, it seems reasonable to conjecture that the concepts of unfolding and transversality should be of use in characterizing various issues arising in the sensitivity analysis of nonlinear programs. Here we shall indicate two different directions to be pursued: 1) constraint qualification conditions; 2) objective function stabilization and examination of the stability of the dual algorithms discussed above.

Transversality and the Kuhn-Tucker Conditions

As an indication of how singularity theory arguments can be employed to study constraint perturbations, let us examine the classical Kuhn-Tucker conditions using transversality arguments.

The Kuhn-Tucker necessary conditions play an important role in the theoretical development of mathematical programming. These conditions were derived from a more general set of conditions, called the Fritz John conditions by assuming that a constraint qualification is in effect. Both the Fritz John and Kuhn-Tucker conditions are necessary for x^* to be an optimal solution of the constrained optimization problem. One of the most widely used constraint qualifications is that the gradients of the binding constraints at the point x^* be linearly independent.

In singularity theory, the concept of a transverse intersection between two manifolds is a cornerstone for structural stability arguments. One definition of a transverse intersection at a point is that no vector is perpendicular to the tangent spaces of both manifolds simultaneously [4]. Since the gradient vector of a manifold at a point will also be the normal vector to the tangent hyperplane at that point, it follows that the gradient vectors of two intersecting manifolds will both be collinear if and only if the intersection is transverse. Furthermore, and more importantly, the Thom Isotropy Theorem [4] states that transverse crossings are structurally stable. This means that small perturbations of the constraints around a Kuhn-Tucker point won't change the geometry of the intersection much. In fact, the original constraint configuration can be recovered by a smooth coordinate change around the point of interest.

Let us consider an example demonstrating the structural instability of a non-transverse crossing. In the example, the following definition of a transverse crossing will be used:

DEFINITION 1. Two manifolds, R and S, embedded in R^n intersect *transversally* if

1) $R \cap S = \phi$ or

2) $\text{codim}(T_x R) + \text{codim}(T_x S) < n$ and

$\text{codim}(T_x R) + \text{codim}(T_x S) = \text{codim } T_x R \cap T_x S$

where T_x is the tangent space at x .

Example (from [7], see Figure 4 for geometry)

Minimize: $f(x_1, x_2) = -x_1$

Subject to: $x_2 - (1-x_1)^3 \leq 0$

$-x_2 \leq 0$

$\nabla f = (-1, 0)$

$\nabla g_1 = (-3(1-x_1)^2, 1)$

$\nabla g_2 = (0, -1)$

at (1,0): $\nabla f(1,0) = (-1, 0)$

$\nabla g_1(1,0) = (0, 1)$ - binding

$\nabla g_2(1,0) = (0, -1)$ - binding

The gradients of the binding constraints are not linearly independent. Checking the Kuhn-Tucker conditions:

$$\nabla f + u_1 \nabla g_1 + u_2 \nabla g_2 = 0$$

$$\begin{bmatrix} -1 \\ 0 \end{bmatrix} + u_1 \begin{bmatrix} 0 \\ 1 \end{bmatrix} + u_2 \begin{bmatrix} 0 \\ -1 \end{bmatrix} = 0$$

$$u_1 \begin{bmatrix} 0 \\ 1 \end{bmatrix} + u_2 \begin{bmatrix} 0 \\ -1 \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \Rightarrow u_1 = 1, u_2 = 0, \text{ inconsistency,}$$

showing that the Kuhn-Tucker conditions don't hold.

Transversality

Both g_1 and g_2 will be embedded in R^3 so

$$T_x g_1(x) = \{(x_1, x_2, x_3) : x_2 - 3(1-x_1)^2 x_1 = x_2 - 3(1-x_1)^2 x_1\}$$

$$T_x g_2(x) = \{(x_1, x_2, x_3) : x_2 = x_2\}$$

at $(1,0)$:

$$T_x g_1(x) = \{(x_1, x_2, x_3) : x_2 = 0\}, \text{ thus } T_x g_1(x) \cap T_x g_2(x) = T_x g_1(x)$$

$$T_x g_2(x) = \{(x_1, x_2, x_3) : x_2 = 0\}$$

$$\text{codim } T_x g_1(x) = 1$$

$$\text{codim } T_x g_2(x) = 1$$

$$\text{codim } (T_x g_1(x) \cap T_x g_2(x)) = 1$$

Thus,

$$\text{codim } T_x g_1(x) + \text{codim } T_x g_2(x) \neq \text{codim } (T_x g_1(x) \cap T_x g_2(x))$$

so the intersection is nontransverse.

If the cubic constraint is perturbed slightly:

$$g_1(y_1, y_2) = y_2 - (1-y_1)^3 + \varepsilon,$$

then $T_x g_1(x) \cap T_x g_2(x)$ at $(1,0)$ will be the empty set, so the intersection is, by definition, transverse. At their point of intersection, $x = (1+\varepsilon, 0)$, so

$$T_x g_1(x) = \{(x_1, x_2, x_3) : x_2 + 3\varepsilon^2 x_1 = 3\varepsilon^2\}$$

$$T_x g_2(x) = \{(x_1, x_2, x_3) : x_2 = 0\}$$

and $T_x g_1(x) \cap T_x g_2(x) = \{(x_1, x_2, x_3) : x_1 = 1\}$ will be a line in R^3 .

Thus,

$$\text{codim } (T_x g(x)) = 1$$

$$\text{codim } (T_x g(x)) = 1$$

$$\text{codim } (T_x g_1(x) \cap T_x g_2(x)) = 2$$

so $\text{codim } (T_x g_1(x)) + \text{codim } (T_x g_2(x)) = \text{codim } (T_x g_1(x) \cap T_x g_2(x))$ and the intersection is transverse.




FIGURE 4 Example of a nontransverse constraint crossing (from Figure 4.5, pg.136, Bazaraa and Shetty, 1979)

The unfolding concept can also be of use in sensitivity analyses in the following manner. As an unparametrized function, the objective function $f(\mathbf{x})$ may be unstable with regard to small perturbations (i.e. the qualitative character of the critical points of f may change as a result of small changes in f). This is clearly a bad situation as far as the credibility of the results obtained from such an optimization are concerned. However, if $\text{codim } f = c$, an unfolding of f involving at least c parameters will be stable relative to all structural perturbations in the sense that if $f(\mathbf{x}) + \mathbf{p}(\mathbf{x})$ is a perturbation of f , then the behavior of $f(\mathbf{x}) + \mathbf{p}(\mathbf{x})$ near its critical points can be captured by varying the parameters in a universal unfolding of f . As already noted, the elements $\mathbf{u}_1(\mathbf{x}), \dots, \mathbf{u}_c(\mathbf{x})$ forming a basis for $\mathfrak{m}_n / \nabla(f)$ constitute a basis for exactly the type of perturbations we need to stabilize f .

Unfolding can also be of use for studying the stability of the dual optimization algorithms, which require the formation of a surrogate objective function using a computational parameter. For example, the augmented Lagrangian method mentioned above requires the use of a parameter ρ . These parameterized functions can be studied to learn what types of objective functions and constraints may lead to surrogate objective

functions which are structurally unstable, and which may behave badly as the computational parameter is varied.

These ideas can be illustrated by considering the standard linear programming problem. In a sense, a linear objective function is the linearization of a general nonlinear $f(x)$, since no real world process even generates a completely linear potential.

DEFINITION 2. f is *structurally stable* if, for sufficiently small smooth perturbation functions p , the critical points of f and $f + p$ remain within the same neighborhood and have the same type (max, min, saddle, etc.).

Consider the linear program:

$$\begin{aligned} \text{maximize: } & f(x) = c^t x \\ \text{subject to: } & Ax \leq 0 \\ & x \geq 0 \end{aligned}$$

Note that the Hessian matrix of f will be identically zero for all x , so that a linear program has a maximum only by virtue of the constraints.

If a small linear perturbation is added to the objective function:

$$\begin{aligned} \text{maximize: } & f(x) = (c^t + \varepsilon^t)x, \varepsilon_i \ll 1 \\ \text{subject to: } & Ax \leq 0 \quad i=1,2,\dots,n \\ & x \geq 0 \end{aligned}$$

the isoclines of the objective function on the x hyperplane might shift so that the set of isoclines leaves the feasible region at a completely different extreme point of the convex hull of constraints. Thus the linear programming problem is not even stable with respect to linear perturbations.

In contrast, it is known that Morse (i.e. quadratic) extrema are the only structurally stable types for nonparameterized functions, although for parameterized functions the situation is different. Similarly, since adding a small perturbation to a Morse function does not drastically change the location of the unconstrained extremum, the location of the constrained extremum also shouldn't change too much, since the constrained extremum usually occurs where the constraints are tangent to the isoclines of $f(x)$.

As a final note, the computational implications of the above discussion are not by any means as dire as might seem. While the general nonlinear programming problem is computationally difficult, numerical methods for quadratic programs, both constrained and unconstrained, are well developed. In fact, since Morse functions are the only structurally stable types of smooth unparametrized functions, a case could be made for transforming even nonquadratic nonlinear programs into quadratic form using the diffeomorphic coordinate changes guaranteed by singularity theory. Thus, a quadratic program represents, in a certain sense, the canonical problem for mathematical programming.

Acknowledgement

Much of the work represented here arose during the course of numerous conversations with J. Kempf who, in particular, is responsible for the example of the use of the gradient projection method.

REFERENCES

1. Arnol'd, V.I. (1981) *Singularity Theory*, Cambridge University Press, Cambridge.
2. Lu, Y.C. (1976) *Singularity Theory and an Introduction to Catastrophe Theory*, Springer, New York.
3. Gibson, G. (1979) *Singular Points of Smooth Mappings*, Pitman, London.
4. Poston, T. and I. Stewart. (1978) *Catastrophe Theory and Its Applications*, Pitman, London.
5. Milnor, J. (1973) *Morse Theory*, Princeton U. Press, Princeton.
6. Golubitsky, M. (1978) *An Introduction to Catastrophe Theory and its Applications*, SIAM Review, 20:352-387.
7. Bazaraa, M.S. and C.M. Shetty. (1979) *Nonlinear Programming*, John Wiley and Sons, New York, 560 pp.
8. Pavelle, R., M. Rothstein, and J. Fitch (1980) *Computer algebra*, Scientific American, 245, C, pp. 136-153.
9. Casti J. (1984) *System Similarities and Natural Laws*, WP-84-1, International Institute for Applied Systems Analysis, Laxenburg, Austria.