

A SURVEY OF NON-LINEAR OPTIMIZATION TECHNIQUES

W. H. SWANN

Wilmslow, England

Optimization means the provision of a set of numerical parameter values which will give the best fit of an equation, or series of equations, to a set of data. For simple systems this can be done by differentiating the equations with respect to each parameter in turn, setting the set of partial differential equations to zero, and solving this set of simultaneous equations (as for example in linear regression). In more complicated cases, however, it may be impossible to differentiate the equations, or very difficultly soluble non-linear equations may result. Many numerical optimization techniques to overcome these difficulties have been developed in the last ten years, and this review explains the logical basis of most of them, without going into the detail of computational procedures.

The methods fall naturally into two classes – direct search methods, in which only values of the function to be minimized (or maximized) are used – and gradient methods, which also use derivatives of the function. The author considers all the accepted methods in each class, although warning that gradient methods should not be used unless the analytical differentiation of the function to be minimized is possible.

If the solution is constrained, that is, certain values of the parameters are regarded as impossible or certain relations between the parameter values must be obeyed, the problem is more difficult. The second part of the review considers methods which have been proposed for the solution of constrained optimization problems.

Introduction

The advent of the digital computer has given a considerable impetus to the study of numerical methods for determining the maximum or minimum of a given function. As a result very many algorithms have been proposed for solving the problem and this survey considers some of those which have been found useful. The paper is intended to be an introduction to the field of non-linear optimization and consequently emphasis will be placed on the ideas on which the algorithms are based rather than on the mathematical and programming details.

The review begins with a statement of the general problem and then considers the unconstrained case, discussing both the direct search and gradient types of method. Techniques used to deal with constraints are then described and finally a few general comments regarding choice of method are made.

Statement of the problem

The problem with which this review is concerned is that of determining the values of a set of parameters x_1, x_2, \dots, x_n , called the independent variables of the problem, which correspond to a minimum of a given

objective function $f(\mathbf{x})$, $\mathbf{x} \equiv (x_1, x_2, \dots, x_n)'$, where the parameters are subject to the m inequality constraints

$$c_i(\mathbf{x}) \geq 0, \quad i = 1, 2, \dots, m,$$

and the s equality constraints

$$e_j(\mathbf{x}) = 0, \quad j = 1, 2, \dots, s \leq n.$$

This is a completely general optimization problem since

$$\begin{aligned} \text{maximum } \{f(\mathbf{x})\} \text{ w.r.t. } \mathbf{x} &= \\ &= \text{minimum } \{-f(\mathbf{x})\} \text{ w.r.t. } \mathbf{x}. \end{aligned}$$

and all constraints can be reduced to the above form.

1. Unconstrained optimization

Many of the methods used for constrained optimization deal with the constraints by converting the

problem in some way into an unconstrained one, and hence it is appropriate to begin the review by considering methods for solving the unconstrained optimization problem.

1.1. Classical approach

Analytically a stationary point of a function $f(\mathbf{x})$ is defined to be one where all of the first partial derivatives of the function with respect to the independent variables are zero, i.e.,

$$\frac{\partial f}{\partial x_i} = 0, \quad i = 1, 2, \dots, n.$$

This stationary point is a minimum if the principal minors of the matrix of second partial derivatives are all positive, i.e.,

$$D_i > 0, \quad i = 1, 2, \dots, n,$$

where

$$D_i = \begin{vmatrix} \frac{\partial^2 f}{\partial x_1^2} & \frac{\partial^2 f}{\partial x_1 \partial x_2} & \dots & \frac{\partial^2 f}{\partial x_1 \partial x_i} \\ \vdots & & & \vdots \\ \frac{\partial^2 f}{\partial x_i \partial x_1} & \dots & \dots & \frac{\partial^2 f}{\partial x_i^2} \end{vmatrix}.$$

Hence the problem could be tackled by differentiating the objective function with respect to each of the variables in turn and equating to zero, which would yield n equations in n unknowns to be solved for the stationary points. However, it may not always be possible to obtain the required derivatives analytically, and even when it is the resulting equations will in most cases be non-linear, and the problem of solving them is no easier than the original optimization problem. Consequently many numerical optimization techniques have been developed, and some of these will now be considered.

1.2. Iterative methods

All numerical optimization techniques except tabulation methods are iterative and starting from an initial approximation \mathbf{x}^0 to the minimum they proceed by defining a sequence of points $\{\mathbf{x}^i\}$, $i = 1, 2, \dots$

in such a way that

$$f(\mathbf{x}^{i+1}) < f(\mathbf{x}^i).$$

This series of improved approximations $\{\mathbf{x}^i\}$ may be considered to be generated by the general iterative equation

$$\mathbf{x}^{i+1} = \mathbf{x}^i + h^i \mathbf{d}^i, \quad (1)$$

where h^i is a positive constant and \mathbf{d}^i is an n -dimensional direction vector evaluated at the i th iteration. The vector \mathbf{d}^i determines the direction to be taken from the i th point \mathbf{x}^i and the magnitude of $h^i \mathbf{d}^i$ determines the size of the step in that direction.

There are many methods in the literature for determining the vector \mathbf{d}^i and they can be divided into two natural classifications — direct search methods and gradient methods. Direct search methods rely solely on values of the objective function; gradient methods use in addition to function values, values of the first and possibly higher order partial derivatives of the function.

1.2.1. Direct search methods

There are many useful methods of the direct search type and it is convenient to further subdivide them into three sub-classes: tabulation methods, linear methods and sequential methods.

(a) Tabulation methods

Tabulation methods assume that the minimum \mathbf{x}^* lies within the region

$$\mathbf{l} \leq \mathbf{x}^* \leq \mathbf{u},$$

where the bounds \mathbf{l} and \mathbf{u} are known. The function is evaluated at the nodes of a grid covering the region of search and the node corresponding to the smallest function value is taken as the minimum. If the range $u_i - l_i$ of variable x_i , $i = 1, 2, \dots, n$, is divided into r_i equal sub-intervals, then the function must be calculated at

$$(r_1 + 1)(r_2 + 1) \dots (r_n + 1)$$

points. Clearly this strategy is very inefficient and it is not recommended.

Random search methods may also be regarded as forms of tabulation. The function is evaluated at points chosen at random from the region of search, with again that point corresponding to the smallest function value taken as the minimum. This too is a very inefficient procedure and is not recommended.

(b) Linear methods

Linear methods are those which use a set of direction vectors during the search which is directed according to the results of explorations along these directions. Some of the methods use the same set of directions throughout the search; others attempt to define new directions along which faster progress may be expected.

(i) Alternating variable method

A first intuitive attempt at a linear direct search optimizing routine might well consist of minimizing along each co-ordinate axis in turn, a procedure which is known as the Alternating Variable Method. The current best point moves parallel to each axis in turn, changing direction when a minimum in the direction being searched is reached, so that if the contours of the objective function are hyperspherical, the minimum will be located after at most n linear searches, starting from the given approximation. This situation is illustrated for $n=2$ in fig. 1 where x^0 is the initial estimate for the minimum x^* which lies at the centre of the concentric circles which are contours of constants function value. x^* is located after two linear searches.




Fig. 1. Alternating variable method for a function of two variables having circular contours of constant f (see text).

However, in general there will be interaction between the variables causing elongation of the contours in some direction, and unless this direction is parallel to one of the co-ordinate axes the search will oscillate




Fig. 2. Alternating variable method for a function of two variables having elliptical contours.

along a slightly inclined valley along the local principal axis of the surface, each step tending to become smaller than the previous one. This case is shown in fig. 2. Hence although very simple the method can prove extremely inefficient, and the inefficiency becomes more pronounced as the number of variables is increased.

(ii) Method of Hooke and Jeeves

Obviously a method which aligns a direction along the principal axis of the contours would be desirable and the method due to Hooke and Jeeves [13] tries to achieve this. The method consists of a combination of exploratory moves and pattern moves: the former seek to locate the direction of any valleys in the surface and the latter attempt to progress down any such valleys.

In an exploratory move each variable is considered in turn and a step δ_i is taken from the current point in the co-ordinate direction x_i . If this results in a decrease in the function the step is successful, the new point becomes the current point and the variable x_{i+1} is considered. Otherwise the step is a failure and is retracted, the sign of δ_i is reversed and a new step taken in the direction x_i (i.e., in the opposite sense). Again if it is successful the new point becomes the current point; otherwise the current point is unaltered. In either event the variable x_{i+1} is then explored in the same manner. This procedure continues until all n variables have been explored and the current point at the end of this search will generally be called a base point.

A pattern move is a step from the current base point, that step having both the magnitude and direction of the line joining the previous base point to the present one.

The method begins by considering the initial approximation as a starting base and making an exploratory move from it. If this exploration fails to produce a direction to search, i.e., if all steps taken in the move are failures, then the starting point is either reasonably close to the minimum or in a sloping valley whose sides are too steep to allow the direction of the valley to be determined using the present step sizes δ_i . In either case the remedy is to reduce the steps δ_i and carry out another exploratory move. If, however, the exploratory move is successful the point reached becomes the new base and a pattern move followed by an exploratory move is made to try to improve the pattern direction. The current function value is then compared with that at the base and if it is less then it becomes the new base and the search continues with a pattern move followed by an exploratory move. When a pattern move followed by an exploratory move fails to improve the function, all steps from the last base are retracted, the base is considered as a starting base and the search recommences from there. Convergence is assumed when the step sizes δ_i have been reduced below some pre-assigned limits. Fig. 3 shows how the method progresses on the function of fig. 2. Starting




Fig. 3. Progress of Hooke and Jeeves' method on a function with elliptical contours.

from x^0 the first exploratory move produces a base point x^1 from where a pattern move is made to x^2 and an exploration to x^3 . The function value at this point is less than that at the base x^4 and so a further pattern move is made to x^4 and an exploratory move to x^5 . Again a pattern step may be taken, giving x^6 , but the result of the ensuing exploration is x^4 which

is an inferior point to x^5 , the present base. Hence all steps from x^5 are retracted and a new exploration made about x^5 , but as can be seen all steps will fail and so the step sizes must be reduced.

Fig. 3 shows how the pattern direction is turned to lie along the principal axis of the contours resulting in much faster progress towards the optimum than was possible using the Alternating Variable Method. The method has been found to be reliable and robust in practice.

(iii) Rosenbrock's method

The method devised by Rosenbrock [20] uses throughout the search a set of n mutually orthonormal directions which are orientated after each iteration so that one of them lies along the direction of total progress made in the iteration.

Suppose we start with a set of n mutually orthonormal directions ξ_i^0 , $i = 1, 2, \dots, n$, usually taken as the co-ordinate directions, and a set of n associated step-sizes δ_i , $i = 1, 2, \dots, n$. A step is defined to be a success if it produces a function value not greater than the value at the point from which the step was taken.

Starting from the initial approximation a step δ_1 is taken along the direction ξ_1^0 , and if it is successful δ_1^0 is multiplied by $\alpha > 1$ so that when the direction is next examined a larger step will be taken. If a failure is recorded the step is retracted and δ_1 is multiplied by β where $-1 < \beta < 0$ so that the next search in direction ξ_1^0 will be made with a smaller step in the opposite sense. In either event the process then goes on to consider the direction ξ_2^0 in the same way and so on until ξ_n^0 has been explored, whereupon the procedure is repeated and continues until a success followed by a failure has been recorded for each direction at some time during the iteration. This occurrence marks the completion of a stage, and before commencing a new one the orthonormal direction vectors are recomputed. ξ_1^1 is given by $(x^1 - x^0)$ normalised, where x^0 and x^1 are the first and last points respectively of the stage just completed, and the remaining vectors ξ_i^1 , $i = 2, 3, \dots$ are calculated from a knowledge of ξ_i^0 , the distances moved along these directions in the previous stage, and ξ_1^1 . The re-orthonormalisation is achieved by using the Gram-Schmidt process.

Rosenbrock suggests values of 3 and -0.5 for the step adjustment parameters α and β , respectively.

The effect of the repeated orthonormalisation is to align ξ_1 along the direction of fastest progress, ξ_2 , along the best direction that can be found normal to ξ_1 and so on, and this overcomes the oscillatory pattern produced by the Alternating Variable Method. The algorithm has been used very widely and has proved to be very reliable; it is particularly useful for locating an early approximation for the minimum.

(iv) The D.S.C. method

The method due to Davies, Swann and Campey, described by Swann [22], also uses a set of orthonormal directions and re-orientates them after each stage, but adopts a different search strategy.

As in Rosenbrock's method n mutually orthonormal direction vectors are chosen, again usually the co-ordinate directions, but in this case a linear minimization is carried out along each one in turn. This linear search is achieved by taking steps along the direction until a bracket on the minimum is obtained, whereupon a quadratic interpolation is used to refine the estimate for the minimum. When each of the directions has been explored once in this manner, new direction vectors are chosen, again taking the direction of total progress during the iteration as the first direction and using the Gram-Schmidt process to determine the others. Those directions in which no progress was made are retained for the next iteration and are excluded from the orthonormalisation. When the distance moved during an iteration is less than the step size δ used in the linear search, δ is reduced; convergence is assumed when δ is less than some pre-set limit. Fig. 4 depicts how the search would proceed on the function of figs. 2 and 3.

The method has generally been found to be more




Fig. 4. Progress of the D.S.C. method on a function of two variables.

efficient than both that of Rosenbrock and that of Hooke and Jeeves.

(v) Powell's method

The method of Powell [16] is one which is based upon conjugate directions and which is quadratically convergent, i.e., it guarantees to locate the minimum of a quadratic objective function of n variables in n iterations. Since most objective functions can be well approximated by quadratics in the neighbourhood of the minimum, this is generally a desirable property.

If the quadratic function to be minimized is

$$f(x) = \frac{1}{2} x' G x + a' x + k,$$

then the directions ξ_i, ξ_j are said to be mutually conjugate with respect to G if

$$\xi_i' G \xi_j = 0.$$

Conjugate directions possess the useful property that the minimum of the function can be located by searching along each of them once only.

The method described by Powell starts with n linearly independent directions and generates conjugate directions by defining a new direction vector after each iteration and replacing one of the current vectors by it. The new direction is again the vector of total progress in the iteration and is added to the end of the list of directions while the first of that list is deleted. This process results in a list of n mutually conjugate directions after n iterations and therefore the exact minimum of a quadratic may be located. For non-quadratic functions the procedure is continued beyond n iterations until during a stage each variable is altered by less than one-tenth of the accuracy required in that variable. Powell does suggest a more stringent alternative to this, but the above criterion has usually proved adequate in ensuring that the minimum is indeed located.

The basic procedure can, however, lead to linearly dependent directions, and to prevent this Powell has modified the algorithm and introduced a criterion to decide if the newly defined vector should be included in the list of directions and if so which vector it should replace.

The method has been used on many problems and has generally proved to be very efficient, with final

convergence especially rapid in the region of the minimum where the function can be well approximated by a quadratic. The modification necessary to ensure that the directions do not become dependent can destroy the quadratic convergence of the method if a recently introduced direction is replaced, and it has been found that on occasion the method fails to replace any direction and the search reduces to an alternating variable procedure.

(c) Sequential methods

Sequential methods are those methods which use some form of geometric configuration to explore the objective function, an approach which originated with *evolutionary operation*, first proposed by Box [1].

(i) The simplex method

A regular simplex in n dimensions is $n + 1$ mutually equidistant points so that in two dimensions it consists of an equilateral triangle, whilst in three dimensions it forms a tetrahedron. It possesses the useful property that a new simplex can be set up on any face of a given simplex by the addition of only one new point and Spendley, Hext and Himsworth [21] made use of this when devising their Simplex Method of optimization.

The search begins by setting up a regular simplex in the space of the independent variables and evaluating the function at each of the $n + 1$ vertices. The vertex corresponding to the greatest function value is then replaced by its reflection in the hyperplane of the remaining points, forming a new simplex. The function is evaluated at the new vertex and the process continued.

There is one exceptional case to be considered. If the vertex with greatest function value is the one most recently introduced, then application of the basic procedure will cause the search to oscillate between the two simplexes. To overcome this difficulty, whenever this situation occurs, the vertex in the new simplex with second largest function value is reflected instead of the largest.

The progress of the search is illustrated in fig. 5 in which the numbers denote the order in which the vertices were introduced. It will be noticed that in the simplex 8-10-11, vertex 8 was reflected rather than vertex 11 as this would have given vertex 9 again. Because vertex 12 is near to the minimum the series of




Fig. 5. Progress of the simplex method on a function of two variables.

simplexes revolves about it, a characteristic feature of the Simplex search. Hence if one vertex of the simplex remains unaltered for a number of consecutive iterations, the size of the figure is diminished by reducing the distances of the remaining vertices from this vertex and recommencing the search.

(ii) The simplex method of Nelder and Mead

The basic simplex routine has been improved by Nelder and Mead [15] who propose that the simplex be altered both in size and configuration by expansion and contraction steps.

The original procedure of reflecting the worst vertex is followed until the new vertex corresponds to either the smallest or largest function value in the new simplex. If the reflection has produced a new best point, then the chosen direction appears to be a profitable one and the simplex is expanded in that direction. The better of the reflected and expanded points is selected as the new vertex. If the reflection has produced a new worst point, then the simplex is contracted back towards the centroid. If the contraction produces a point which would not be the worst of the new simplex then it is chosen as the new vertex; otherwise a new simplex is set up by halving the distance of the vertices from the best vertex. Convergence is assumed when the standard deviation of the function values at the $n + 1$ vertices is less than some specified value.

The basic reflection, expansion and contraction steps are illustrated in fig. 6 for a function of two variables where X_r is the reflected point, X_e is the




Fig. 6. Nelder and Mead's reflection, expansion and contraction steps for a function of two variables.

the expanded point and X_c is the contracted point.

The method is robust and has been used to solve a number of problems efficiently. However, the efficiency diminishes as the number of variables increases, and for n greater than five the procedure is rather inefficient.

1.2.2. Gradient methods

Gradient methods are those methods which use values of the partial derivatives of the function with respect to the independent variables in addition to values of the function itself.

(a) Steepest descent methods

The direction of fastest progress or "steepest descent" at any given point is the direction whose components are proportional to the first partial derivatives of the function at that point. Cauchy [6] is credited with the first application of the steepest descent direction to optimization and many variations of using the direction have subsequently been proposed.

A basic variation would be to define as a search direction d^i the normalised gradient vector at the current point:

$$d^i = \frac{\left(\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \dots, \frac{\partial f}{\partial x_n} \right)}{\left[\sum_{j=1}^n \left(\frac{\partial f}{\partial x_j} \right)^2 \right]^{\frac{1}{2}}}$$

This direction is used with a specified step size h^i to obtain a new trial point from the iterative equation

$$x^{i+1} = x^i + h^i d^i.$$

This procedure is repeated until a step is tried which does not cause a function improvement which indi-




Fig. 7. Progress of steepest descent for a function of two variables.

cates that h^i should be reduced. Fig. 7 shows typical progress for a function of two variables in which the step must be reduced before progress can be made from x^2 .

One of the more often used variants of the method of steepest descent searches along the direction d^i as defined above for the minimum before calculating d^{i+1} . Successive directions are orthogonal and the search is therefore similar to the alternating variable process and is usually very inefficient.

(i) Newton's method

In an attempt to improve the convergence of gradient methods consider the Taylor series expansion of $f(x)$ about the minimum x^* where $x = x^* + \delta$.

If g is the vector of first order partial derivatives of the function and G the matrix of second order partial derivatives, i.e.,

$$g_i = \frac{\partial f}{\partial x_i}, \quad G_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}, \quad i, j = 1, 2, \dots, n,$$

then

$$f(x) = f(x^* + \delta) \approx f(x^*) + \delta'g + \frac{1}{2}\delta'G\delta. \quad (2)$$

At the minimum all the first derivatives are equal to zero, so if (2) is exact then the gradient vector at the current point x must satisfy

$$g = G\delta. \quad (3)$$

The minimum x^* is therefore obtained from x by the move $x^* = x - \delta$ where δ is the solution of the simultaneous linear equations (3), i.e.,

$$x^* = x - G^{-1}g. \quad (4)$$

Hence if the function is a quadratic, the minimum can be located by applying (4) with \mathbf{g} evaluated at the current point \mathbf{x} and \mathbf{G} evaluated at the minimum and the method is quadratically convergent. The minimum is not known but for a quadratic \mathbf{G} is constant and can be evaluated at the current point.

When the function is not a quadratic an iterative approach must be adopted and in Newton's method \mathbf{g} and \mathbf{G} are calculated at the current point \mathbf{x}^i and a further approximation to the minimum is obtained by using

$$\mathbf{x}^{i+1} = \mathbf{x}^i - (\mathbf{G}^i)^{-1} \mathbf{g}^i.$$

This method has two drawbacks, however. Firstly the computation of the matrix of second derivatives and its inversion is likely to prove very time-consuming. Secondly progress towards the minimum is only ensured if \mathbf{G} is positive definite. Hence although Newton's method is efficient in the neighbourhood of the minimum where the function approximates a quadratic and the matrix of second derivatives is positive definite, away from the minimum it is likely to progress only slowly and it may even diverge.

(ii) Davidon's method

The method due originally to Davidon [7] and subsequently refined by Fletcher and Powell [11] is one which begins as steepest descent, gradually accumulates information concerning the curvature of the objective function and uses this information to obtain improved search directions, and converges on the minimum using Newton's method, but does so without resorting to the calculation of second derivatives.

The basic iteration is defined as

$$\mathbf{x}^{i+1} = \mathbf{x}^i - h^i \mathbf{H}^i \mathbf{g}^i,$$

where \mathbf{g}^i is the gradient vector evaluated at \mathbf{x}^i and \mathbf{H}^i is the i th approximation to the inverse of the matrix of second derivatives. The initial approximation to \mathbf{G}^{-1} , i.e., \mathbf{H}^0 , is arbitrary provided that it is positive definite, and the unit matrix is usually chosen so that the first iteration proceeds as steepest descent. The step h^i is chosen so that \mathbf{x}^{i+1} is the minimum along the direction $-\mathbf{H}^i \mathbf{g}^i$, i.e., a linear search is carried out along this direction. After locating \mathbf{x}^{i+1} the estimate

for \mathbf{G}^{-1} is improved according to

$$\mathbf{H}^{i+1} = \mathbf{H}^i + \mathbf{A}^i + \mathbf{B}^i,$$

where \mathbf{A}^i and \mathbf{B}^i are matrices calculated from the progress made during the last iteration and the change this caused in the gradient vector. One of these terms ensures that the matrix \mathbf{H} remains positive definite, while the other ensures that $\mathbf{H} \rightarrow \mathbf{G}^{-1}$ so that for an n -dimensional quadratic $\mathbf{H}^n = \mathbf{G}^{-1}$ and the minimum can be located with a Newton step. Hence the method is quadratically convergent.

The convergence criteria are based on the full Newton step $-\mathbf{H}^i \mathbf{g}^i$ and the actual step taken $-h^i \mathbf{H}^i \mathbf{g}^i$ and the search is terminated if either (i) every component of both these vectors is less than some specific value or (ii) the lengths of both vectors are less than a specified size.

The Davidon method has been used widely on a large number of problems and has proved to be very efficient. It is generally considered to be the most powerful of the currently available optimization techniques, although there is now considerable interest in finding improved methods of updating the matrix \mathbf{H} .

2. Constrained optimization

The classical method of solving the constrained optimization problem:

$$\begin{aligned} &\text{minimize } f(\mathbf{x}) \\ &\text{subject to } c_i(\mathbf{x}) \geq 0 \quad i = 1, 2, \dots, m, \\ &\text{and } e_j(\mathbf{x}) = 0 \quad j = 1, 2, \dots, s \leq n \end{aligned}$$

uses Lagrangian multipliers to convert the problem into an unconstrained one. In doing so, however, it produces a saddle-point problem which is more difficult to solve than the original constrained problem, and hence the usefulness of this approach is very limited.

The feasible region

A point \mathbf{X} in the parameter space at which all of the constraints are satisfied, i.e.,

$$c_i(\mathbf{X}) \geq 0, \text{ all } i; \quad e_j(\mathbf{X}) = 0, \text{ all } j,$$




Fig. 8. Feasible region ABCDE defined by six constraints.

is said to be feasible and the entire collection of such points constitutes the feasible region. All other points are non-feasible and constitute the non-feasible region.

In fig. 8 the constraints are shaded on the non-feasible side so that ABCDE defines the boundary of the feasible region and all points inside that boundary are feasible. As fig. 8 demonstrates the constraints may exclude the optimum M of the objective function from the feasible region, and in such cases the constrained optimum x^* will generally lie on the boundary of the feasible region.

In most iterative methods for constrained optimization an initial feasible point must be provided, and in problems involving a number of non-linear constraints it may be difficult to find such a point. A useful method of obtaining a feasible point from a non-feasible one is to minimize the sum of the constraint violations:

$$\min C = - \sum_i C_i(x) + \sum_{j=1}^s e_j^2(x),$$

where the optimization is unconstrained and the first summation runs over only those of the m inequality constraints which are currently violated. A minimum of zero indicates that a feasible point has been located, but failure to converge to such a minimum does not indicate that a feasible point does not exist, merely that the search has failed to locate one.

2.1. Transformations

Before considering methods of handling constraints, it is worth noting that constraints can often be eliminated by transforming the variables of the problem.

For example, if the independent variable x is subject to the constraint

$$x \geq k \quad \text{or} \quad l \leq x \leq u,$$

where k, l, u are constants, then the problem can be converted into an unconstrained one in the variable y by writing

$$x = k + y^2 \quad \text{or} \quad x = l + (u-l) \sin^2 y.$$

It sometimes occurs in practical problems that the only constraints on the problem are of the above forms and hence the technique can prove very useful. A paper by Box [3] discusses transformations and gives some further examples.

2.2. Intuitive approach

Most of the techniques for unconstrained optimization consist of a sequence of linear searches so that an initial attempt to extend them to handle constraints might be to arrange that, whenever a constraint is violated, return to the last feasible point and recommence the search with a reduced step, continuing until a feasible minimum is located even if that minimum lies on the constraint.

This approach can be very efficient for univariate problems, but in a multivariate search it can lead to premature termination of the minimization. For example using either Powell's method or the D.S.C. method to minimize the function of fig. 9, the direction ξ_1 might be searched first and, applying the above rule, the point x^1 chosen as the minimum in that direction. No progress can be made along the direction ξ_2 , however, so no new directions can be defined and the search will be terminated with x^1 indicated as the required minimum.




Fig. 9. Difficulties with the intuitive approach.

This example would seem to indicate that what is required is a method which follows any constraints encountered during the search, and a number of such techniques have been devised.

2.2.1. Elimination of variables

One method which follows constraints uses the effective constraints to eliminate variables. For example, if $f(x)$ is to be minimized subject to

$$\sum_{j=1}^n c_j x_j \geq 0,$$

then the optimization is carried out unconstrained until the constraint is violated, whereupon x_1 is written in terms of x_2, \dots, x_n , i.e.,

$$x_1 = -\frac{1}{c_1} \sum_{j=2}^n c_j x_j,$$

and the optimization is continued in the remaining $n-1$ variables. This has the effect of restricting the search to move along the constraint as required and has the added advantage that it reduces the dimensions of the problem.

2.2.2. Riding the constraint

However, it is not always easy or possible to use the violated constraint to express one of the variables in terms of the others, and when this is the case the constraint may be followed using the technique of "riding the constraint" due to Roberts and Lyvers [17].

Again the minimization is carried out unconstrained until a constraint is violated, whereupon the current point is advanced to the constraint boundary either by taking repeatedly smaller steps or by some form of interpolation. If the constraint is

$$c(x) \geq 0,$$

then along the boundary, i.e., $c(x) = 0$, $dc = 0$ since the constraint represents a contour of constant value. Therefore

$$dc = \frac{c}{x_1} dx_1 + \frac{c}{x_2} dx_2 + \dots + \frac{c}{x_n} dx_n = 0. \quad (5)$$

In this relationship the partial derivatives are evaluated at the current point. If to move along the constraint a step $dx = (dx_1, dx_2, \dots, dx_n)$ must be taken, then $n-1$ of the increments dx_i can be specified by the search routine and the remaining one is determined by eq. (5). Immediately a new constraint is violated the method switches over to ride it, provided that the function continues to improve.

Non-linear constraints can be followed quite effectively by either of these methods, but both suffer from the drawback that they assume that the minimum lies on a constraint. Generally, this will be true, but there are cases when the minimum is unconstrained and even if good progress can be achieved by following a constraint for some time it may eventually be preferable or necessary to leave the constraint. For example in fig. 9 fast progress can be made from x^1 by moving along $c_1 = 0$, but the minimum x^* is unconstrained. Neither elimination of variables nor riding the constraint makes provision for leaving a constraint.

2.2.3. Hemstitching

A method which can move along a constraint but which does not assume that the minimum lies on a constraint is the method of "mathematical hemstitching", also due to Roberts and Lyvers [17]. In this method, immediately a constraint is violated the search is returned to the feasible region by taking a step orthogonal to the constraint. Hence, if the search is continually moving into the non-feasible region the path of progress will be repeatedly crossing the constraint boundary and can be said to be hemstitching along it (fig. 10). If two or more constraints are violated a return direction is set up using a weighted sum of the constraint gradients.

The main difficulty with this method is that there is no guarantee that the point in the feasible region to which the search returns is an improvement on the best point obtained before leaving the feasible space, for example in fig. 10, the function value at x^{i+1} is greater than that at x^i . Consequently progress can be very slow or even non-existent and the process may degenerate into a random search.

2.2.4. Penalty functions

A different approach to the problem of constrained minimization is to weight the objective function so




Fig. 10. Hemstitching along a constraint using the method of steepest descent.

that non-feasible points are unattractive to the search.

A possible weighting when minimizing a function $f(x)$ subject to

$$c_i(x) \geq 0, \quad i = 1, 2, \dots, m,$$

would be achieved by forming

$$F(x) = f(x) + \sum_{i=1}^m k_i c_i^2 H(c_i),$$

where $H(c_i)$ is the Heaviside unit step function of argument c_i and the k_i are positive weights. The function $F(x)$ is then minimized without taking further account of the constraints. This penalty function has the effect that in the feasible region the true function is minimized, but when the non-feasible region is entered the function is increased by a weighted sum of the squared constraint violations. It can be shown that under certain conditions the unconstrained minimum of $F(x)$ tends to the constrained minimum of $f(x)$ as the weights k_i tend to infinity. Any unconstrained optimization routine may be used to minimize $F(x)$ with the weights k_i being successively modified as the search proceeds. A convenient scheme for the calculation of the k_i is given by Leitmann [14].

The method works reasonably well, but creates steep valleys and discontinuous derivatives at the constraint boundary and these features are often difficult to overcome, particularly when using gradient methods.

2.2.5. Rosenbrock's method

A more complex method of imposing a penalty on

the objective function was described by Rosenbrock [20] who considered constraints of the type

$$l_i \leq x_i \leq u_i, \quad i = 1, 2, \dots, m,$$

where the implicit variables x_{n+1}, \dots, x_m are functions of the explicit variables x_1, \dots, x_n and the upper and lower bounds $u_i, l_i, i = 1, 2, \dots, m$, are either constant or functions of the explicit variables.

Boundary zones

$$l_i \leq x_i \leq l_i + 10^{-4} (u_i - l_i), \quad i = 1, 2, \dots, m,$$

$$u_i \geq x_i \geq u_i - 10^{-4} (u_i - l_i), \quad i = 1, 2, \dots, m,$$

are introduced at the edges of the range of each variable, explicit and implicit, and the objective function is modified only within these boundary areas so that in the interior of the feasible region the function is unaltered.

During the optimization any step which causes a constraint to be violated is considered to be a failure. If a step enters a boundary zone then the function is modified so that the innermost limit of the zone the function is unaltered, and at the outside limit, i.e., on the constraint, the function takes the best value obtained before the boundary zone was entered. This modification is achieved by defining the fractional depth of penetration λ of the boundary zone as

$$\lambda = \frac{\text{amount by which boundary zone is entered}}{\text{width of boundary zone}}$$

and then minimizing not $f(x)$ but a function $\phi(x)$

where

$$\phi(x) = f - (f - f^*)(3\lambda - 4\lambda^2 + 2\lambda^3)$$

and f^* is the best value of the function for which no boundary zone was entered. Clearly $\phi(x)$ has the required properties at both extremes of the boundary zone, i.e., at $\lambda = 0$ and $\lambda = 1$, so that a turning point is introduced within that zone. Fig. 11 illustrates the case for a function f of a single variable x , considering the upper limit only.




Fig. 11. Rosenbrock's function modification for a function of one variable.

This method of applying constraints has worked well in conjunction with Rosenbrock's optimization technique, but it does not appear to be very effective when applied to other search methods.

2.2.6. The created response surface technique

The Created Response Surface Technique, devised by Carroll [5] and developed by Fiacco and McCormick [9], also modifies the objective function, but does so in such a way as to convert the constrained problem into a series of unconstrained ones. This is achieved by creating a sequence of new response surfaces, each one of which moves the optimum nearer to the true optimum of the objective function. The function modification involved is similar to that of Rosenbrock in that it becomes increasingly severe as a constraint is approached, whilst in the interior of the feasible region the function is relatively unaltered.

If the function $f(x)$ is to be minimized subject to

the inequality constraints $c_i(x) \geq 0$, $i = 1, 2, \dots, m$, then a new response surface can be created by defining

$$\phi(x) = f(x) + k \sum_{i=1}^m \frac{w_i}{c_i(x)}, \quad (6)$$

where k and w_i are positive constants. With this formulation, as one of the constants approaches its limit of zero the corresponding penalty increases so that in the limit $\phi(x)$ tends to infinity as $c_i(x)$ tends to zero and the surface is asymptotic to the constraint. Hence the created function $\phi(x)$ has a feasible minimum. In the interior of the region, of course, the penalty has little effect on the function. The case for a function f of one variable x subject to the constraint $c(x) \geq 0$ is shown in fig. 12. Because the created minimum must always be feasible no further account need be




Fig. 12. A created response surface for a function of one variable.

taken of the constraints and the problem of locating the minimum of $\phi(x)$ can be considered an unconstrained one. Obviously the minimum of the created surface might be remote from the constrained minimum of $f(x)$ depending upon the choice of the weights k and w_i , so a number of surfaces are formed, each successive surface bearing a less severe penalty so that the created minimum approaches the desired minimum.

Hence given an initial feasible approximation x^0 to the solution, the point x^1 is located where

$$x^1 = \min_x \{ \phi(x, k_1) \},$$

$$\phi(x, k_1) = f(x) + k_1 \sum_{i=1}^m \frac{w_i}{c_i(x)},$$

and

$$k_1 > 0.$$

x^1 is an improved feasible approximation to the constrained minimum. Starting from x^1 the point x^2 is found where

$$x^2 = \min_x \{\phi(x, k_2)\}$$

and

$$0 < k_2 < k_1.$$

Continuing in this manner a sequence of points $\{x^r\}$, $r = 1, 2, \dots$, is generated that approximates to the minima of $\{\phi(x, k_r)\}$ where $\{k_r\}$ is a strictly monotonic decreasing sequence and $k_r \rightarrow 0$ as $r \rightarrow \infty$.

Fig. 13 depicts how the search might proceed for the problem

$$\begin{aligned} \text{minimize } f(x) &= -x \\ \text{subject to } c_1(x) &= x \geq 0 \\ c_2(x) &= 1-x \geq 0 \end{aligned}$$




Fig. 13. Created response surface technique for a function of one variable.

The initial value for the weight k_1 and the factor by which it is reduced to obtain subsequent values can be guessed intuitively and unless very bad guesses are made convergence will be reasonable. Too small an initial value causes the first surface $\phi(x, k_1)$ to approximate $f(x)$ too closely, the smoothing property of the technique is lost and convergence to the first optimum

can be very slow. Too large an initial value causes the first surface to be only slightly dependent on the objective function so that the created surface is forced well into the interior of the feasible region and the usefulness of a good initial approximation to the constrained minimum would therefore be lost. Fiacco and McCormick have suggested two methods for computing k_1 , but since these methods make use of derivatives of the function and the penalty they can be rather time consuming, and since they are not completely reliable a guessed value may be used.

Fiacco and McCormick have reported that the amount of work required to locate the optimum is relatively independent of the rate of decrease of k . The faster the reduction the less surfaces to be optimized, but each minimum is more difficult to locate. Slower reduction necessitates more surfaces, each of which is easier to optimize.

There remains the problems of choosing the weights w_i and Fiacco and McCormick suggest that they are best set equal to unity. The method in this form certainly works, but experience has shown that better results can be obtained using more realistic values, although no satisfactory algorithm has yet been proposed for their automatic determination.

This form of penalty function is widely used, although a number of variations of it have been proposed, for example Box, Davies and Swann [4] square $c_i(x)$ in eq. (6) and take the summation over only those constraints that have been violated at any time during the search.

Fiacco and McCormick [10] have extended the method to apply to equality constraints of the form

$$e_j(x) = 0, \quad j = 1, 2, \dots, s,$$

by forming the response surface

$$\begin{aligned} \phi(x, k_r) &= f(x) + k_r \sum_{i=1}^m \left(\frac{w_i}{c_i(x)} \right) \\ &+ k_r^{\frac{1}{2}} \sum_{j=1}^s [e_j(x)]^2. \end{aligned}$$

Note that in this case the initial estimate x^0 to the solution does not need to satisfy the equality constraints.

Fiacco and McCormick have shown that under certain conditions the sequence of minima will converge to the constrained minimum of the true objective function. The method has been used successfully on a wide variety of problems, including many for which it is not possible to prove that convergence is assured, e.g., problems with non-convex feasible regions.

2.2.7. The complex method

A different approach was adopted by Box [2] who modified the Simplex method of Spendley et al., so that it recognizes constraints. This constrained version uses $k > n$ points to form what is termed a complex and performs a Simplex-type minimization until a constraint is encountered whereupon it rolls the complex along the constraint in the direction of decreasing f , and it can leave the constraint at a later stage if this is desirable. This search is called the *complex method*.

Suppose the problem is to minimize $f(\mathbf{x})$ subject to the constraints

$$l_i \leq x_i \leq u_i, \quad i = 1, 2, \dots, m \geq n.$$

The variables $x_{n+1}, x_{n+2}, \dots, x_m$ are called the implicit variables and are functions of the explicit variables x_1, x_2, \dots, x_n . Then given an initial feasible point \mathbf{x}^0 further $k-1$ points \mathbf{x}^p are chosen to form the complex. A trial point is generated according to

$$x_i^p = l_i + r_i^p(u_i - l_i), \quad i = 1, 2, \dots, n,$$

where the parameters r_i^p are pseudo-random rectangularly distributed deviates in the interval (0, 1). A point generated in this manner must necessarily satisfy the explicit constraints, but it may violate one or more of the implicit ones. If the trial point is feasible it is accepted as a vertex of the complex; otherwise it is retracted halfway back towards the centroid of the points already chosen including \mathbf{x}^0 , this move being repeated until ultimately the point is feasible.

The function is evaluated at each of the k vertices, and as in the simplex method the vertex having greatest function value is replaced. Initially the vertex is replaced by its over-reflection in the centroid of the remaining vertices where by over-reflection is meant the point on the produced line joining the rejected point to the centroid, but α times as far from the centroid as the mirror image of the reflected point,

with the reflection coefficient α greater than unity. If this new point satisfies all the constraints the function is evaluated there and the process repeated, unless the new point is also the worst vertex in the new complex in which case repeated moves halfway back towards the centroid are made until the new vertex is not the worst. If the new point does not satisfy some explicit constraint, the corresponding variable is reset just inside the appropriate boundary. If the new point violates an implicit constraint, repeated moves halfway back towards the centroid are made until a feasible point is obtained. Box originally suggested that convergence should be assumed when five consecutive equal function evaluations have been made; alternatively the terminating criterion could be based on the standard deviation of the function values at the vertices as recommended by Nelder and Mead.

It is necessary to use more than $n+1$ points as vertices to prevent the complex collapsing prematurely into a sub-space and Box suggests $k = 2n$, although such a value will be too large for problems with as many as ten variables. A value in excess of unity for the expansion coefficient α will enlarge the complex and compensate for contractions towards the centroid and will accelerate progress when the initial approximation is remote from the minimum; Box recommends a value of 1.3.

This method has been used successfully to solve a number of problems and is a useful method for constrained optimization.

2.2.8. Projection methods

None of the methods described thus far makes any distinction between linear and non-linear constraints and there is at present considerable interest in methods using projection matrices since such methods can take account of the special case of linear constraints.

If \mathbf{N} is the matrix whose columns are the gradients of the active constraints

$$c_i(\mathbf{x}) = 0, \quad i = 1, 2, \dots, m,$$

i.e.,

$$\mathbf{N} = \begin{bmatrix} c_{11} & \dots & c_{m1} \\ c_{12} & & \cdot \\ \cdot & & \cdot \\ \cdot & & \cdot \\ c_{1n} & & c_{mn} \end{bmatrix}$$

with

$$\sum_{j=1}^n c_{ij}^2 = 1, \quad i = 1, 2, \dots, m,$$

then the projection matrix

$$P = I - N(N'N)^{-1}N' \quad (7)$$

projects any vector into the intersection of the constraints if they are linear, or into the intersection of their tangent hyperplanes if they are non-linear.

(a) Rosen's gradient projection method

Probably the earliest of the projection methods was that due to Rosen [18,19] which applies projection matrices to the method of steepest descent.

The minimization is performed using a search vector the direction of steepest descent $g = \nabla f$ until an inequality constraint is violated whereupon the current point is moved to the constraint boundary by interpolation. The constraint may now be considered to be an equality one and the gradient direction is projected into its tangent hyperplane, i.e., the new search direction is Pg . Fig. 14 illustrates the case for




Fig. 14. Gradient projection for a function of two variables subject to a non-linear constraint.

a function of two variables. If the constraint is linear the search direction is feasible since it lies along the constraint, but if it is non-linear and convex as in fig. 14 then a step along Pg will produce a non-feasible point and an interpolation routine must be used to return to the feasible region down some return direction. If a number of constraints are active then the

gradient direction is projected into the intersection of their tangent hyperplanes. If at any stage the gradient direction is feasible with respect to any of the constraints previously violated, then that constraint is deleted from N and hence from P so that the search may leave that constraint.

This method, whilst a useful means of handling constraints, suffers from the usual criticisms of steepest descent in that it takes no account of the curvature of the objective function and can therefore exhibit the usual oscillatory behaviour resulting in slow progress.

2.2.9. Linear constraints with Davidon's method

The unconstrained gradient technique due to Davidon is generally considerably more efficient than the method of steepest descent and Holdfarb and Lapidus [12] have extended the method to take into account linear constraints using projection operators without destroying its power.

The optimization begins in the feasible region with all of the inequality and equality constraints satisfied. It will be recalled that Davidon's method proceeds according to

$$x^{i+1} = x^i - h^i H^i g^i,$$

where g^i is the gradient vector evaluated at x^i , H^i is the i th approximation to the inverse of the matrix of second derivatives and h^i is the distance moved along $-H^i g^i$ so that x^{i+1} is a minimum in that direction. Goldfarb and Lapidus suggest that the rank of H be reduced so that all search directions are projected to lie in the subspace formed by the intersection of the equality constraints. The method then proceeds using the Davidon iteration including the updating of H which preserves the projection nature of H . As in Rosen's method, when an inequality constraint is violated it is considered as an equality and the rank of H further reduced so that the search can proceed along the constraint boundary, and if any such inequality constraints are subsequently found to be no longer active the rank of H is increased again to allow the search to leave the constraint.

The application of projection operators to Davidon's method in this way produces a very powerful algorithm for solving problems with linear constraints that converges for a quadratic function subject to p

active constraints in $n-p$ iterations. It has also been used successfully in conjunction with the Created Response Surface Technique to deal with non-linear constraints, thereby producing a useful method for the general optimization problem.

2.2.10. Davies' hemstitching method

Davies [8] has proposed extending the use of projection matrices to handle non-linear constraints with a method based on the philosophy of the hemstitching algorithm of Roberts and Lyvers.

In the original hemstitching process when a constraint is violated a step is taken perpendicular to the constraint to return to the feasible region, and the main criticism of the method is that it can degenerate to a random search because there is no guarantee that the new point will be an improvement. Davies overcomes this objection by adopting the following approach.

Basically the method proceeds in the same manner as proposed by Goldfarb and Lapidus until a non-linear constraint is violated whereupon the search direction is projected into the tangent hyperplane to the constraint. For a convex constraint this will produce a direction \mathbf{d}^i , steps along which will be non-feasible as shown in fig. 15, so Davies suggests a return




Fig. 15. Davies' hemstitching method for a non-linear constraint.

direction based on the magnitudes and gradients of the active constraints. In fig. 15, \mathbf{r}^i represents this return direction \mathbf{d}^i being the search direction $-\mathbf{H}^i \mathbf{g}^i$, and AB is the intersection of the plane normal to the steepest descent direction \mathbf{g}^i at \mathbf{x}^i and the plane defined by the search direction \mathbf{d}^i and the return direction \mathbf{r}^i . This line represents a limit beyond which any

return point will not improve the function value if the objective function is concave. Hence if \mathbf{u}^i is the point reached by a step along \mathbf{d}^i from \mathbf{x}^i and \mathbf{v}^i is the intersection of AB with \mathbf{r}^i then the point \mathbf{w}^i on the constraint boundary can be found by interpolation and a new step taken from \mathbf{w}^i in the direction \mathbf{d}^i . This process is continued until the minimum \mathbf{x}^{i+1} along this constraint using the direction \mathbf{d}^i is located, and then the matrix \mathbf{H} can be updated and a new search direction defined.

The method has been tried on a number of realistic test examples and the results obtained show considerable promise.

2.3. Conclusions

There is at present no universal optimization routine which will solve any given problem more efficiently than any other method and it should always be borne in mind that there are some problems for which those algorithms usually considered to be inefficient may prove very useful.

The initial choice lies between direct search and gradient methods, and gradient methods should only be used if the derivatives of the function can be obtained analytically — differentiation by differences is not recommended — and of all gradient methods currently available that proposed by Davidon is undoubtedly the most efficient. If analytic derivatives are not available then a direct search method should be used and of these Powell's algorithm is probably the most effective, although it can under certain conditions prove very inefficient if it fails to choose any new directions. In this event the D.S.C. method is a good alternative. The method of Nelder and Mead is useful on problems of less than about five variables and the methods of Rosenbrock, and Hooke and Jeeves, have also been used successfully on many practical problems.

Of the methods available for handling constraints, the complex method, Rosenbrock's method applied to his own search routine and Carroll's Created Response Surface Technique applied to a number of search procedures, both gradient and direct search, have all been used extensively, and generally Carroll's approach has usually proved to be the most effective. Wherever possible, however, a method should take advantage of linear constraints, and if function and constraint gradients are available analytically, limited

experience suggests that either Goldfarb and Lapidus' method used in conjunction with Carroll's technique for non-linear constraints, or Davies' hemstitching method should be used.

References

- [1] G.E.P.Box, Evolutionary operation: a method for increasing industrial productivity, *Appl. Statistics* 2 (1957) 81.
- [2] M.J.Box, A new method of constrained optimization and a comparison with other methods, *The Computer Journal* 8 (1965) 42.
- [3] M.J.Box, A comparison of several current optimization methods, and the use of transformations in constrained problems, *The Computer Journal* 9 (1966) 66.
- [4] M.J.Box, D.Davies and W.H.Swann, Non-linear optimization techniques, I.C.I. Monograph No. 5 (Oliver and Boyd, Edinburgh, 1969).
- [5] C.W.Carroll, The created response surface technique for optimizing non-linear restrained systems, *Operations Res.* 9 (1961) 169.
- [6] A.L.Cauchy, Méthode générale pour la résolution des systèmes d'équations simultanées, *Compt. Rend. Acad. Sci. Paris* 25 (1847) 536.
- [7] W.C.Davidon, Variable metric method for minimization, A.E.C. Research and Development Report ANL-5990 (Rev.) (1959).
- [8] D.Davies, The use of Davidon's method in non-linear programming, I.C.I. Ltd., Management Services Report, MSDH/68/110 (1968).
- [9] A.V.Fiacco and G.P.McCormick, The sequential unconstrained minimization technique for non-linear programming, a primal-dual method, *Management Sci.* 10 (1964) 360.
- [10] A.V.Fiacco and G.P.McCormick, Extensions of SUMT for non-linear programming: equality constraints and extrapolation, *Management Sci.* 12 (1966) 816.
- [11] R.Fletcher and M.J.D.Powell, A rapidly convergent descent method for minimization, *The Computer Journal* 6 (1963) 163.
- [12] D.Goldfarb and L.Lapidus, Conjugate gradient method for non-linear programming problems with linear constraints, I. and E.C. Fundamentals 7 (1968) 142.
- [13] R.Hooke and T.A.Jeeves, Direct search solution of numerical and statistical problems, *J. A. C. M.* 8 (1961) 212.
- [14] G.Leitmann (editor), Optimization Techniques with Applications to Aerospace Systems (Academic Press, New York, 1962).
- [15] J.A.Nelder and R.Mead, A simplex method for function minimization, *The Computer Journal* 7 (1965) 308.
- [16] M.J.D.Powell, An efficient method of finding the minimum of a function of several variables without calculating derivatives, *The Computer Journal* 7 (1964) 155.
- [17] S.M.Roberts and H.I.Lyvers, The gradient method in process control, *Ind. Eng. Chem.* 53 (1961) 877.
- [18] J.B.Rosen, The gradient projection method for non-linear programming. Part I. Linear constraints, *J. Soc. Indust. Appl. Math.* 8 (1960) 181.
- [19] J.B.Rosen, The gradient projection method for non-linear programming. Part II. Non-linear constraints, *J. Soc. Indust. Appl. Math.* 9 (1961) 514.
- [20] H.H.Rosenbrock, An automatic method for finding the greatest or least value of a function, *The Computer Journal* 3 (1960) 175.
- [21] W.Spendley, G.R.Hext and F.Himsworth, Sequential application of simplex designs in optimisation and evolutionary operation, *Technometrics* 4 (1962) 441.
- [22] W.H.Swann, Report on the development of a new direct search method of optimization, I.C.I. Ltd., Central Instrument Laboratory Research Note 64/3 (1964).