

Algorithm for Linear Programming

Anand Mohan Sinha¹ & Kumar Mukesh²

¹Department Of Mathematics, J. P. University, Chapra, Bihar- 841301, (INDIA.)

²Research Scholar, Department Of Mathematics, J. P. University, Chapra, Bihar-84130,1(INDIA.)

Abstract: In this paper we will describe the “An Algorithm for Linear Programming”.

Key words and phrases: Active Set, inactive Set & karush kuhn tucker (KKT) conditions.

AMS Subject Classification 2010: 90C05

I. Introduction

This algorithm has three main components: Selecting which components to include in the next update of \hat{x} , deciding how far to move in that direction, and when a constraint restricts a move to decide on a pivot column.

Most of the work occurs in making transformations of variables and in checking the KKT conditions [8-11], so some extra work elsewhere is likely to be helpful if it reduces on average the number of these actions. Given a direction, d in which to move, each $\delta \hat{x}_i$ that is part of that direction contributes to an inner product that defines how much to change the $\hat{x}_i, i > n$ for a given move μd . Each nonzero component of d adds to this cost so it makes sense not to include directions for which there is a fairly high probability that the desired direction will change as a result of future transformations. Finally there is a cost of deciding just how far it is possible to move in a given direction.

Each nonzero d_i is equal to $-\hat{g}_i$. Other strategies we have tried have not been as successful. The choice of components to include in d is simply based on the size of the $|\hat{g}_i|$. We include the largest one, and any that are “close”. The initial choice of μ is simply the smallest value which just reaches a bound for an associated x_i . If a value of 0 is computed, then that variable is made active and at least for the time being is not considered for inclusion in d . This situation changes if this \hat{g}_i at some later point is a sufficiently large violation of the KKT conditions.

Using the inner products, the individual $\hat{x}_i, i > n$ are checked to find the first, if any that is reached. If so, a pivot is selected and transformations performed to convert the \hat{x}_i from being a general constraint into being a variable, and the whole process starts over. Otherwise the variable that is reached is made active, and if there is more than 1 variable left in d the inner products are updated to account for the variable just removed, a new μ is computed and the search repeated.

Of course if the μ is unbounded and no constraint is encountered, then the minimum is unbounded. And if the KKT conditions are all satisfied and all constraints are active we have found the minimum.

For the pivot column we select the column with the coefficient of largest magnitude, after adding a penalty for pivot choices which would have the new constraint immediately fail the KKT test.

When the code is not making much progress, factors are attached to constraints which get smaller the more times a constraint goes active or inactive. Variables with small factors are less likely to be made active when more than one constraint limits the move, and are less likely to be made inactive as a result of checking the KKT conditions. Because rounding errors can make progress impossible on problems which should be solved, constraint violations are ignored when inner products are quite small.

It may happen that equality has no pivot choice which is larger than might be expected from computational noise. Such equality is dependent on others. If it also has a large residual then the equalities are inconsistent and the computations are terminated with an indication that this is the case. Otherwise the equality is redundant and is moved into an area where it is ignored.

This same dependency problem can occur with the gradient. That is, all directions which are available have $|\hat{g}_i|$ values which are so small that their sign is in question. In this case an option in the code lets one choose

between accepting the current point as the solution, and continuing the computation by attempting to minimize x_i for the smallest value of i which allows such freedom. This process is then continued until no freedom is left thus, insuring the solution point returned lies on a vertex, i.e. the solution will not be an interior point of some hyper plane.

II. Formulation of problem

For this purpose we consider a problem. Our problem of interest is

$$\text{Subject to bounds on } Cx - b \text{ and } X, \text{ minimize } c^T X, \quad (1)$$

Where X is a real n -vector C is a real matrix with m_c rows and n columns, and b is a real m_c -vector. The problem is linear in the sense that with exact arithmetic it can be solved exactly in a finite number of steps. When upper and lower bounds are equal the condition is treated as an equality constraint. It is convenient to treat c^T as a row of C with an unbounded value, and this is what the algorithm does.

Unlike the simplex algorithm, our algorithm makes no effort to stay on a vertex and does not introduce slack variables to deal with linear inequalities.

The algorithm internally works in a transformed coordinate system. The variables in the current coordinate system are identified by \hat{x} . A given \hat{x}_k may correspond to an original x_j or to a transformed variable generated when making a constraint active. In the internal co-ordinates $\hat{a}_i^T \hat{x} - \hat{b}_i$, may refer to some other row in the original matrix or to one of the original variables. We (at least attempt to) attach \hat{a} to variables in the internal coordinate system. \hat{C} Over-writes C in memory. Initially the first n cells of \hat{x} are identical to x , and the following m_c cells are used for the initial value of $Cx - b$.

The first n cells of \hat{x} one can think of as variables of which the last α are active and currently have fixed values. The first $n - \alpha$ of these variables are free to change in value, although at a given time only a subset of them do so. A method of this type is said to employ an active set strategy.^[1-7]

Variable transformations are accomplished using elementary transformations from the right. This is like the standard LU factorization.

If we assume the constraint row is \hat{a}_i^T , and that the pivot element is $a_{i,p}$. The transformations involve multiplying \hat{C} from the right with a matrix of the form $I + e_p t^T$, where e_p is 0 except for a 1 in row p , $t_k = 0$ for $i \geq p$ and for $i < p$, $t_j = -\hat{a}_{i,j} / \hat{a}_{i,p}$. At the same time the vector \hat{x} must be updated by multiplying it by the inverse of this matrix, $I - e_p t^T$. A given row of physical storage may thus contain several logical rows. The right most part of the row will always corresponds to original rows in the constraint matrix, but when such a row is active other parts would corresponds to variables which have been converted to general constraints. All but the leftmost part of the row corresponds to active constraints. The value associated with the right hand side is now corresponds to the value for the inactive variable now stored in the leftmost part of the row.

The algorithm insures that variables are always logically ordered from lowest indexed to highest (or left to right) as: inactive simple variables, active variables and active inequalities, factored equalities, equalities on bounds, and variables to be ignored. Finally in the code, variables with an index $> n$ correspond to inactive variables which have become general constraints and then other inactive inequalities.

When thinking about the various transformations we sometimes find it useful to imagine that we start with additional rows in our coefficient matrix of the form $\hat{I}\hat{X} = x$. As we apply transformations this changes to $\hat{X} = x$, and for every row here that is not a row of the identity matrix, there is a row someplace else that has been factored. We recommend thinking in these terms if you find something confusing about the transformations.

Let us mention a few of the terms used later

Active Set: The set of variables which are on a bound and not free to change in value.

Inactive Set: The set of variables which are free to change in value.

A subset of this change on each iteration.

General Constraint: A constraint involving a linear combination of what are currently the variables.

III. Karush Kuhn Tucker (KKT) Conditions

The current gradient, \hat{g} , gives the direction in which the objective is increasing most rapidly. If \hat{g}_i is positive we would like to decrease \hat{x}_i in order to reduce the objective. If x_i is setting on a lower bound such a decrease is not possible and \hat{g}_i is said to satisfy the KKT conditions with respect to that condition.

The Lagrange multipliers or dual variables at a solution satisfy $[IC^T]\lambda = g$, (where the comes from the conditions on the first n x_i) and since we can multiply this equation on the left by the transpose of the transformations we are making, we can also write $[IE^T]\lambda = \hat{g}$. The λ_i 's associated with the inactive constraints are all 0. Thus the conditions on the Lagrange multipliers and on \hat{g} can be written

Condition	Active State	Satisfies KKT	Violates KKT
=	Satisfied	always	never
\geq	On a lower bound	$\hat{g}_i \geq 0$	$\hat{g}_i \leq 0$
\leq	On an upper bound	$\hat{g}_i \leq 0$	$\hat{g}_i \geq 0$

IV. Getting Started

If the user doesn't specify an initial point, $x = 0$ is used. The first step in the starting procedure is to ensure that all the initial variables satisfy their bounds. Any that do not are immediately set to the nearest bound, and the $\hat{x}_i, i > n$ are updated to account for the change.

Any variables that do not affect the value of any \hat{x}_i for $i > n$ are permuted to a high index where they can be neglected from this point on. Such variables are set to the value of the bound that gives the minimal value of the objective. If the bounds are equal for a variable that variable is immediately converted to equality on that variable and the variable is permuted to the right (a higher index) where it is ignored from this point on. Empty rows are likewise permuted to a location where they no longer need to be examined.

By default rows and columns are scaled to balance the sizes of the rows and columns appearing in E . This simplifies the later selection of good pivots, and the tests that are made for removing a constraint from the active set.

The next step is to process the equality constraints. After the factorization, the \hat{x}_i for $i > n$ are updated to account for the change necessary to satisfy the equality constraint. Note that in this process some of the original variables have turned into general constraints, and that these constraints may be violated as a result of satisfying the equality conditions. The variables and columns that have become active as a result are ignored except when done, where a back solve is necessary in order to compute the associated Lagrange multipliers. Thus when done with the equalities there are fewer variables, and no equality constraints. Unfortunately at this time we may have a number of constraints that are not satisfied. However constraints on variables \hat{x}_i for $i \leq n$ are always satisfied except for some drift which can occur when having problems with rounding errors or in rare cases when having trouble getting feasible.

V. Getting Feasible

When constraints are violated, the algorithm substitutes its own objective in place of that from the user. The objective in this case is to minimize the L_1 norm of the violated constraints. The gradient in this case is made up of additions and subtractions of the rows of E corresponding to the violated constraints. The row will be added in if the upper bound is violated, and subtracted if the lower bound is violated. The algorithm is very similar to that for the linear programming case, except that checks must be made for previously violated constraint becoming feasible and when they do to adjust the gradient for the change. There are no infeasible constraints during this phase as what were infeasible constraints initially have been converted into being part of the objective during this phase. The algorithm used is quite similar to that described for the linear programming case below.

References:

- [1] A. V. Fiacco and G.P. McCormick (1968), "Nonlinear Programming: Sequential Unconstrained Minimization Techniques", Wiley.
- [2] A. Foresgren, P. E. Gill and M. H. Wright (2002), "Interior methods for nonlinear optimization", *SIAM Review* 44, 525–597.
- [3] R. W. Freund, F. Jarre and S. Schaible (1996), "On self-concordant barrier functions for conic hulls and fractional programming", *Math. Program*, 74, 237-246.
- [4] K. R. Frisch (1955), "The logarithmic potential method of convex programming" *Memorandum of May 13, University Institute of Economics, Oslo, Norway*.
- [5] D. M. Gay, M. L. Overton and M. H. Wright (1998), "A primaldual interior method for nonconvex nonlinear programming" *In Advances in Nonlinear Programming (Y. Yuan, ed.), Kluwer, pp. 31-56*.
- [6] P. E. Gill, W. Murray, M. A. Saunders, J. A. Tomlin and M. H. Wright (1986), "On projected Newton barrier methods for linear programming and an equivalence to Karmarkar's projective method" *Math. Program.* 36, 183-209.
- [7] M. X. Goemans (1997), "Semidefinite programming in combinatorial optimization" *Math. Program*, 79, 143-161.
- [8] C. C. Gonzaga (1989), "An algorithm for solving linear programming problems in $O(n^3L)$ operations. In *Progress in Mathematical Programming: Interior Point and Related Methods (N. Megiddo, ed.), Springer, New York, pp. 1-28*.
- [9] N. I. M. Gould, D. Orban and P. L. Toint (2005), "Numerical methods for large-scale nonlinear optimization" *In Acta Numerica, Vol. 14, Cambridge University Press, pp. 299-361*.
- [10] C. Gueret, C. Prins and M. Sevaux (2002), "Applications of Optimization with XpressMP, Dash Optimization" *Translated and revised by Susanne Heipcke*.
- [11] O. Guler (1996), "Barrier functions in interior-point methods", *Math. Oper. Res.* 21, 860-885.