

A Newton Method for Linear Programming¹

O. L. MANGASARIAN²

Abstract. A fast Newton method is proposed for solving linear programs with a very large ($\approx 10^6$) number of constraints and a moderate ($\approx 10^2$) number of variables. Such linear programs occur in data mining and machine learning. The proposed method is based on the apparently overlooked fact that the dual of an asymptotic exterior penalty formulation of a linear program provides an exact least 2-norm solution to the dual of the linear program for finite values of the penalty parameter but not for the primal linear program. Solving the dual problem for a finite value of the penalty parameter yields an exact least 2-norm solution to the dual, but not a primal solution unless the parameter approaches zero. However, the exact least 2-norm solution to the dual problem can be used to generate an accurate primal solution if $m \geq n$ and the primal solution is unique. Utilizing these facts, a fast globally convergent finitely terminating Newton method is proposed. A simple prototype of the method is given in eleven lines of MATLAB code. Encouraging computational results are presented such as the solution of a linear program with two million constraints that could not be solved by CPLEX 6.5 on the same machine.

Key Words. Linear programming, Newton method, least norm solution, exterior penalty.

¹The research described in Data Mining Institute Report 02-02 was supported by National Science Foundation Grants CCR-9729842 and CCR-0138308, by Air Force Office of Scientific Research Grant F49620-00-1-0085, and by the Microsoft Corporation. I am indebted to Robert R. Meyer for suggesting the MATLAB command SPONES to help generate very large test problems, to Jinho Lim for providing a CPLEX mex file for MATLAB, to Steve Wright for suggesting the use of various CPLEX options, and to my PhD student Michael Thompson for the numerical results presented in Table 3.

²John von Neumann Professor Emeritus of Mathematics and Computer Sciences, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin.

1. Introduction

The method proposed here is motivated by the effectiveness of a finitely terminating Newton method proposed in Ref. 1 for the unconstrained minimization of strongly convex piecewise quadratic functions arising from quadratic programs and utilized successfully for classification problems in Ref. 2. To apply this approach to linear programs, a reasonable choice is the least 2-norm formulation (Refs. 3–6) of a linear program as a strongly convex quadratic program, which gives an exact least 2-norm solution for finite parameter values. This has been done in Ref. 7, where a finite Newton method was proposed but without any computational results and without giving an exact solution to the dual of the least 2-norm solution. In our approach here, we make no assumption about our linear program other than solvability and an implied uniqueness condition, which are explained in detail in Section 2. In Section 3, we state our Newton algorithm with an Armijo stepsize and give its global convergence. In Section 4, we give encouraging comparative test results with CPLEX 6.5 (Ref. 8) on a class of synthetically generated sparse linear programs with as many as two million constraints as well as on six publicly available machine learning classification problems. We also give brief MATLAB codes for a test problem generator as well as a simple version of our Newton solver without an Armijo stepsize. This code is used to obtain all of our numerical results. On five out of seven synthetic test problems, the proposed method was faster than CPLEX by a factor in the range of 2.8 to 34.2. On the remaining two problems, CPLEX ran out of memory. On the six smaller machine learning classification problems, CPLEX was faster but LPN gave higher training and testing set classification correctness.

We note that, under the implied uniqueness assumption of our paper, all that is needed is that the number of constraints m of the primal linear program (2) be no less than the number of variables n of the problem. However the principal effectiveness of the proposed method is to problems where m is much bigger than n as evidenced by the numerical examples of Section 4. Other approaches to linear programs with very large number of constraints are given in Refs. 9–10. Another related approach that uses similar ideas to those presented here is given in Ref. 11, but does not recover the primal solution from the dual as we do.

A word about our notation. All vectors will be column vectors unless transposed to a row vector by a prime superscript. For a vector x in the n -dimensional real space R^n , the plus function x_+ is defined as

$$(x_+)_i = \max\{0, x_i\}, \quad i = 1, \dots, n,$$

while x_* denotes the subgradient of x_+ which is the step function defined as

$$\begin{aligned} (x_*)_i &= 1, & \text{if } x_i > 0, \\ (x_*)_i &= 0, & \text{if } x_i < 0, \\ (x_*)_i &\in [0, 1], & \text{if } x_i = 0, i = 1, \dots, n. \end{aligned}$$

The scalar (inner) product of two vectors x and y in the n -dimensional real space R^n will be denoted by $x'y$. The 2-norm of x will be denoted by $\|x\|$, while $\|x\|_1$ and $\|x\|_\infty$ will denote the 1-norm and ∞ -norm respectively. For a matrix $A \in R^{m \times n}$, A_i is the i th row of A which is a row vector in R^n and $\|A\|$ is the 2-norm of A , $\max_{\|x\|=1} \|Ax\|$. If $S \subset \{1, \dots, m\}$, the A_S is the submatrix of A consisting of rows $A_{i \in S}$. A column vector of ones of arbitrary dimension will be denoted by e and the identity matrix of arbitrary order will be denoted by I . If f is a real-valued function defined on the n -dimensional real space R^n , the gradient of f at x is denoted by $\nabla f(x)$ which is a column vector in R^n and the $n \times n$ matrix of second partial derivatives of f at x is denoted by $\nabla^2 f(x)$. For a piecewise quadratic function, such as

$$f(x) = (1/2)\|(Ax - b)_+\|^2,$$

where $A \in R^{m \times n}$ and $b \in R^m$, the ordinary Hessian does not exist because its gradient, the $n \times 1$ vector

$$\nabla f(x) = A'(Ax - b)_+,$$

is not differentiable. However, one can define its generalized Hessian (Refs. 1, 12, 13), which is the $n \times n$ symmetric positive-semidefinite matrix

$$\partial^2 f(x) = A' \text{diag}(Ax - b)_* A, \tag{1}$$

where $\text{diag}(Ax - b)_*$ denotes an $m \times m$ diagonal matrix with diagonal elements $(A_i x - b_i)_*$, $i = 1, \dots, m$. The generalized Hessian (1) has many of the properties of the regular Hessian (Refs. 1, 12, 13) in relation to $f(x)$. Throughout this work, the notation $:=$ will denote definition and s.t. will stand for “such that”.

2. Equivalence of Primal Exterior Penalty LP to Dual Least 2-Norm LP

We give in this section an apparently overlooked, but implied (Refs. 3, 7, 14), result that a parametric exterior penalty formulation of a linear program for any sufficiently small but finite value of the penalty parameter, provides an exact least 2-norm solution to the dual linear program. We begin

with the primal linear program

$$\min_{x \in R^n} c'x, \quad \text{s.t. } Ax \leq b, \quad (2)$$

where $c \in R^n$, $A \in R^{m \times n}$, $b \in R^m$, and its dual

$$\max_{u \in R^m} -b'u = -\min_{u \in R^m} b'u, \quad \text{s.t. } A'u + c = 0, u \geq 0. \quad (3)$$

We write the parametric exterior penalty formulation of the primal linear program for a fixed positive value of the penalty parameter ϵ as the unconstrained minimization problem

$$\min_{x \in R^n} f(x), \quad (4)$$

where f is the penalty function

$$f(x) := \min_{x \in R^n} \epsilon c'x + (1/2) \|(Ax - b)_+\|^2. \quad (5)$$

The positive penalty parameter ϵ needs to approach zero in order to obtain a solution to the original linear program (Refs. 15, 16) from (4). However, this will not be the case if we look at the least 2-norm formulation (Refs. 3, 5) of the dual linear program (3),

$$-\min_{v \in R^m} (b'v + (\epsilon/2)v'v), \quad \text{s.t. } A'v + c = 0, v \geq 0. \quad (6)$$

That (6) leads to a least 2-norm solution of the dual linear program (3) follows from the fact established in Refs. 3, 5 that, for any positive ϵ such that $\epsilon \in (0, \bar{\epsilon}]$ for some positive $\bar{\epsilon}$, the minimization problem (6) picks, among the elements of the solution set of the dual linear program (3), that which minimizes the perturbation term $v'v/2$. Because the objective function of (6) is strongly convex, its solution \bar{v} is unique. The necessary and sufficient Karush-Kuhn-Tucker optimality conditions for problem (6) are that there exists a $y \in R^n$ such that

$$\epsilon v \geq 0, \quad \epsilon v + b - Ay \geq 0, \quad \epsilon v'(\epsilon v + b - Ay) = 0, \quad A'v + c = 0, \quad (7)$$

or equivalently,

$$\epsilon v = (Ay - b)_+, \quad A'v + c = 0. \quad (8)$$

That is,

$$v = (1/\epsilon)(Ay - b)_+, \quad A'(Ay - b)_+ + \epsilon c = 0. \quad (9)$$

Defining $f(y)$ as in (5), the optimality conditions (9) for the dual least 2-norm solution become

$$v = (1/\epsilon)(Ay - b)_+, \quad \nabla f(y) = A'(Ay - b)_+ + \epsilon c = 0. \quad (10)$$

That is,

$$v = (1/\epsilon)(Ay - b)_+, \quad y \in \arg \min_{y \in R^n} f(y) = \arg \min_{y \in R^n} (1/2)\|(Ay - b)_+\|^2 + \epsilon c'y, \tag{11}$$

which are precisely the necessary and sufficient conditions that y be a minimum solution of the parametric exterior penalty function $f(y)$ for the primal linear program (2), for any positive value of the penalty parameter ϵ . Hence, solving the exterior penalty problem (4) for any positive ϵ provides a solution

$$v = (1/\epsilon)(Ay - b)_+$$

to (6). If in addition $\epsilon \in (0, \bar{\epsilon}]$ for some positive $\bar{\epsilon}$, it follows by Theorem 1 of Ref. 5 that v is a least 2-norm solution to the dual linear program (3). We have thus established the following proposition.

Proposition 2.1. Equivalence of Least 2-Norm Dual to Finite-Parameter Penalty Primal. The unique least 2-norm solution to the dual linear program (3) is given by

$$v = (1/\epsilon)(Ay - b)_+, \tag{12}$$

where y is a solution of the primal penalty problem,

$$\min_{y \in R^n} f(y) = (1/2)\|(Ay - b)_+\|^2 + \epsilon c'y, \tag{13}$$

for any finite value of the penalty parameter $\epsilon \in (0, \bar{\epsilon}]$ for some positive $\bar{\epsilon}$.

We note that the gradient,

$$\nabla f(y) = A'(Ay - b)_+ + \epsilon c, \tag{14}$$

which is Lipschitz continuous with constant $\|A'\| \|A\|$, is not differentiable. However, as stated in the Introduction, a generalized Hessian (Refs. 1, 12, 13) with many of the properties of the ordinary Hessian can be defined, which is the following $n \times n$ symmetric positive-semidefinite matrix:

$$\partial^2 f(y) = A' \text{diag}(Ay - b)_* A, \tag{15}$$

where $\text{diag}(Ay - b)_*$ denotes an $m \times m$ diagonal matrix with diagonal elements $(A_i y - b_i)_*$, $i = 1, \dots, m$. The step function $(\cdot)_*$, is defined in the Introduction and is implemented here with $(0)_* = 0$. The matrix $\partial^2 f(y)$ will be used to generate our Newton direction, or more precisely a modified Newton direction, since the generalized Hessian may be singular in general. In fact, the direction that will be used is the following one:

$$d := -(\delta I + \partial^2 f(y))^{-1} \nabla f(y), \tag{16}$$

where δ is some small positive number. With this direction and a variety of stepsizes (Ref. 17, Example 2.2), we can establish global convergence. A key empirical computational property of this direction appears to be global convergence for a class of linear programs with $m \gg n$ from any starting point without any stepsize at all. The exact least 2-norm solution \bar{v} for the dual linear program (3) is obtained by first solving the primal exterior penalty problem (13) for any $\epsilon \in (0, \bar{\epsilon}]$, and then using (12) to determine the unique least 2-norm dual solution \bar{v} . This exact dual solution can be utilized to generate an exact solution to the original primal linear program (2) by solving the constraints of the linear program (2) corresponding to positive components of \bar{v}_j as equalities, that is,

$$A_j z = b_j, \quad j \in S := \{j \mid \bar{v}_j > 0\}. \quad (17)$$

We note that this system of linear equations must always have a solution z as a consequence of the complementarity condition:

$$A_j z - b_j = 0, \quad \text{for } \bar{v}_j > 0, j = 1, \dots, m.$$

In fact, (17) yields an exact solution of the primal linear program (2) if we make the additional assumption that

$$\text{the submatrix } A_S \text{ has linearly independent columns,} \quad (18)$$

where S is defined in (17). This assumption, which implies that the solution of the linear system (17) is unique, is sufficient but not necessary for generating an exact primal solution as will be shown in Section 3.

We turn now to our algorithmic formulation and its convergence.

3. Linear Programming Newton Algorithm: LPN

Our proposed algorithm consists of solving (13) for an approximate solution y of the linear program (2), computing an exact least 2-norm solution v to the dual linear program (3) from (12), and finally computing an exact solution z to the primal linear program (2) from (17). In order to guarantee global convergence, we utilize an Armijo stepsize (Refs. 18, 19) and need to make the linear independence assumption (18) on the least 2-norm solution of the dual linear program (3). We now state our algorithm.

Algorithm 3.1. LPN: Linear Programming Newton Algorithm. Set the parameter values ϵ , δ and tolerance tol (typically 10^{-3} , 10^{-4} , and 10^{-12} respectively). Start with any $y^0 \in R^n$ [typically $y^0 = \bar{A}'\bar{A} + \epsilon I)^{-1} \bar{A}'\bar{b}$, where \bar{A}

is an arbitrary $n \times n$ subset of A and \bar{b} is the corresponding $n \times 1$ subset of b . For $i = 0, 1, \dots$, do the following steps.

- (I) Compute $y^{i+1} = y^i - \lambda_i(\partial^2 f(y^i) + \delta I)^{-1} \nabla f(y^i) = y^i + \lambda_i d^i$, where the Armijo stepsize $\lambda_i = \max\{1, 1/2, 1/4, \dots\}$ is such that

$$f(y^i) - f(y^i + \lambda_i d^i) \geq -(\lambda_i/4) \nabla f(y^i)' d^i, \tag{19}$$

and d^i is the modified Newton direction,

$$d^i = -(\partial^2 f(y^i) + \delta I)^{-1} \nabla f(y^i). \tag{20}$$

- (II) Stop if $\|y^i - y^{i+1}\| \leq \text{tol}$. Else, set $i = i + 1$ and go to (I).
- (III) Define the least 2-norm dual solution v as

$$v = (1/\epsilon)(Ay^{i+1} - b)_+ \tag{21}$$

and a solution z of the primal linear program by

$$A_j z = b_j, \quad j \in S := \{j | v_j > 0, j = 1, \dots, m\}. \tag{22}$$

We state a convergence result for this algorithm now.

Theorem 3.1. Each accumulation point \bar{y} of the sequence $\{y^i\}$ generated by Algorithm 3.1 solves the exterior penalty problem (4). The corresponding \bar{v} obtained by setting y and \bar{y} in (12) is the exact least 2-norm solution to the dual linear program (3), provided ϵ is sufficiently small. An exact solution \bar{z} to the primal linear program (2) is obtained from (22) by solving for z with $v = \bar{v}$, provided that the submatrix $A_{\bar{S}}$ of A has linearly independent columns, where

$$\bar{S} := \{j | \bar{v}_j > 0, j = 1, \dots, m\}. \tag{23}$$

Proof. Let $\epsilon > 0$. That each accumulation point \bar{y} of the sequence $\{y^i\}$ solves the minimization problem (13) follows from standard results [Ref. 17, Theorem 2.1, Examples 2.1(i) and 2.2(iv)], the fact that the direction choice d^i of (20) satisfies

$$-\nabla f(y^i)' d^i = \nabla f(y^i)' (\delta I + \partial^2 f(y^i))^{-1} \nabla f(y^i) \geq (\delta + \|A' A\|)^{-1} \|\nabla f(y^i)\|^2, \tag{24}$$

and the fact that we are using an Armijo stepsize (19). Now let $\epsilon \in (0, \bar{\epsilon}]$. Then by Proposition 2.1, the corresponding \bar{v} obtained by setting y to \bar{y} in (12) is the exact least 2-norm solution to the dual linear program (3). If the submatrix A_S has linearly independent columns, then the solution \bar{z} of (22) with $v = \bar{v}$ is unique and must be a solution of the primal linear program (2). \square

Remark 3.1. Choice of ϵ . Determining the size of $\bar{\epsilon}$, such that the solution v of the quadratic program (6), for $\epsilon \in (0, \bar{\epsilon}]$, is the least 2-norm solution of the dual problem (3) is not an easy problem theoretically. However, computationally this does not seem to be critical and is effectively addressed as follows. By Ref. 20, Corollary 3.2, if for two successive values of ϵ , $\epsilon^1 > \epsilon^2$, the corresponding solutions of the ϵ -perturbed quadratic programs (6), u^1 and u^2 , are equal, then under certain assumptions, $u = u^1 = u^2$ is the least 2-norm solution of the dual linear program (3). This result is implemented computationally by using an ϵ , which when decreased by a factor of 10 yields the same solution to (6).

We note that the assumptions of Theorem 3.1, and in particular the uniqueness of the solution of (22) with $v = \bar{v}$, imply the uniqueness of the solution to the primal linear program (2). This follows from the fact that the least 2-norm multiplier \bar{v} is a valid optimal multiplier for all optimal solutions of the primal linear program (2) and hence all primal solutions must equal the solution of the uniquely solvable system $A_{\bar{S}}z = b_{\bar{S}}$. We term our condition of linear independence of the columns of $A_{\bar{S}}$ as a strong uniqueness condition, because it implies primal solution uniqueness but is not implied by it. Example 3.1 below has a unique but not strongly unique solution. However, it is still solvable by our proposed Newton method despite the fact that its solution is not strongly unique.

Example 3.1. Consider the primal and dual problems

$$\min_{x \in \mathbb{R}^2} x_1, \quad \text{s.t. } x_1 + x_2 \leq -1, x_1 - x_2 \leq 1, -x_1 \leq 0, \quad (25)$$

$$\max_{0 \leq u \in \mathbb{R}^3} u_1 - u_2, \quad \text{s.t. } u_1 - u_2 + u_3 = 1, -u_1 + u_2 = 0. \quad (26)$$

The unique primal solution is

$$x_1 = 0, \quad x_2 = -1.$$

The dual solution set is

$$\{u \in \mathbb{R}^3 \mid u_3 = 1, u_1 = u_2 \geq 0\},$$

and the least 2-norm dual solution is

$$\bar{u}_1 = \bar{u}_2 = 0, \quad \bar{u}_3 = 1.$$

Thus,

$$A_{\bar{S}} = A_3 = [-1, 0],$$

and the solution $x_1 = 0, x_2 = -1$ is not strongly unique because the columns of A_3 are not linearly independent. However, Algorithm 3.1 solved this

problem in 5 iterations using the MATLAB Code 4.2 with the given defaults and without an Armijo stepsize and yielding $z = [0, -1]'$ and $v = [0, 0, 1]'$.

We note further that Algorithm 3.1 can possibly handle linear programs with even nonunique solutions as evidenced by the following simple example.

Example 3.2. Consider the primal and dual problems

$$\min_{x \in R^2} x_1 + x_2, \quad \text{s.t. } x_1 - x_2 \leq -1, -x_1 \leq 0, -x_2 \leq 0, \tag{27}$$

$$\max_{0 \leq u \in R^3} u_1, \quad \text{s.t. } u_1 + u_2 = 1, u_1 + u_3 = 1. \tag{28}$$

The primal solution set is

$$\{x \in R^2 \mid x_1 + x_2 = 1, x_1 \geq 0, x_2 \geq 0\},$$

while the dual solution set consists of the unique point

$$u_1 = 1, \quad u_2 = u_3 = 0.$$

When Algorithm 3.1 is applied to this example using the MATLAB Code 4.2 with the given default values, the exact primal $z = [1, 0]'$ and dual $v = [1, 0, 0]'$ solutions were obtained in two iterations without using the Armijo stepsize.

We give now a sufficient (but in general not necessary) condition for the sequence $\{y^i\}$ of Algorithm 3.1 to converge.

Corollary 3.1. The sequence $\{y^i\}$ is bounded and hence has an accumulation point and the corresponding sequence $\{v^i\}$ converges to the least 2-norm solution of the dual linear program (3), provided that R^n is contained in the conical hull of the rows of $\begin{bmatrix} c' \\ A \end{bmatrix}$. That is, for each $p \in R^n$,

$$\zeta c' + s' A = p, \quad (\zeta, s) \geq 0, \quad \text{has a solution } (\zeta, s) \in R^{1+n}. \tag{29}$$

Proof. We show prove that the sequence $\{y^i\}$ is bounded, hence it has an accumulation point. Since all accumulation points of the sequence $\{v^i\}$, that correspond to accumulation points of $\{y^i\}$, are all equal to the unique least 2-norm solution \bar{v} of (3), the sequence $\{v^i\}$ must converge to \bar{v} .

We show now that the sequence $\{y^i\}$ is bounded by exhibiting a contradiction if it were unbounded and if assumption (29) holds.

Suppose that $\{y^i\}$ is unbounded. Then, for some subsequence (for simplicity, we drop the subsequence index), we have by Algorithm 3.1 that

$$f(y^i) \leq f(y^0), \quad y^i \rightarrow \infty. \tag{30}$$

Dividing by $\|y^i\|^2$ gives

$$(1/2)\|(Ay^i/\|y^i\| - b/\|y^i\|)_+\|^2 + \epsilon c' y^i/\|y^i\|^2 \leq f(y^0)/\|y^i\|^2. \quad (31)$$

Letting $y^i \rightarrow \infty$ and denoting by \hat{y} an accumulation point of the bounded sequence $y^i/\|y^i\|$, we have that

$$(1/2)\|(A\hat{y})_+\|^2 \leq 0. \quad (32)$$

Hence,

$$A\hat{y} \leq 0, \quad \hat{y} \neq 0. \quad (33)$$

From (31) multiplied by $\|y^i\|$ and noting that its first term is nonnegative, we have that

$$\epsilon c' y^i/\|y^i\| \leq f(y^0)/\|y^i\|, \quad (34)$$

from which, upon letting $y^i \rightarrow \infty$, we get

$$c' \hat{y} \leq 0. \quad (35)$$

Combining (33) and (35) gives

$$A\hat{y} \leq 0, \quad c' \hat{y} \leq 0, \quad \hat{y} \neq 0. \quad (36)$$

However, this contradicts assumption (29) if we set $p = \hat{y}$ in (29) as follows:

$$0 = \zeta c' \hat{y} + s' A\hat{y} - \hat{y}' \hat{y} \leq -\hat{y}' \hat{y} < 0. \quad (37)$$

□

Remark 3.2. We note that condition (29) of Corollary 3.1 is equivalent to the primal linear program having a bounded level set, which is similar to the assumptions made in Ref. 7. To see this, we have by the Farkas theorem (Ref. 21, Theorem 2.4.6) that (29) is equivalent to

$$Ax \leq 0, \quad c'x \leq 0, \quad p'x > 0 \text{ has no solution } x \in R^n, \quad \text{for each } p \in R^n. \quad (38)$$

This in turn is equivalent to

$$Ax \leq 0, \quad c'x \leq 0, \quad x \neq 0 \text{ has no solution } x \in R^n. \quad (39)$$

This is equivalent to the boundedness of the level set $\{x | Ax \leq b, c'x \leq \alpha\}$ of the linear program (2) for any real number α .

We proceed now to the numerical testing of our algorithm.

4. Numerical Tests

The objectives of the preliminary numerical tests presented here are to display the simplicity of the proposed LPN algorithm, its competitiveness

with a state-of-the-art linear programming code for special types of problems, and its ability to provide exact answers for a class of linear programs with a very large number of constraints and such that $m \gg n$. Typically, $m \geq 10n$. We demonstrate also the effectiveness of the LPN algorithm by testing it and comparing it with CPLEX on standard classification test problems: four from the University of California Machine Learning Repository (Ref. 22) and two publicly available datasets (Ref. 23).

4.1. Very Large Synthetic Datasets. We first test LPN on very large synthetically generated test problems. To display the simplicity of LPN, we give two MATLAB m-file codes below. The first, `lpgen`, is a linear programming test problem generator. The second, `lpnewt1`, is an implementation of Algorithm 3.1 without the apparently unnecessary Armijo stepsize for the class of problems tested here. A sufficient well-conditioned property that eliminates the Armijo stepsize requirement (Ref. 1, Theorem 3.6) apparently is not needed in all the problems tested here. Both m-files are given below as Code 4.1 and Code 4.2, and are available at: www.cs.wisc.edu/math-prog/olvi.

The test problem generator `lpgen` generates a random constraint matrix A for given m , n , and density d . The elements of A are uniformly distributed between -50 and $+50$. A primal random solution x , with elements in $[-10, 10]$ approximately half of which are zeros, and a dual solution u , with elements in $[0, 10]$ approximately $3n$ of which are positive, are first specified. These solutions are then used to generate an objective function cost vector c and a right-hand side vector b for the linear program (2). The number $3n$ of positive dual variables is motivated by support vector machine applications (Refs. 24–26) where the number of positive multipliers corresponding to the support vectors is often a few multiples of the dimensionality n of the input space.

To test LPN, we solved a number of linear programs generated by the `lpgen` Code 4.1 on LOCOP2, a 400 Mhz Pentium II machine with a maximum of 2 Gigabytes of memory running Windows NT Server 4.0 and compared it with CPLEX 6.5 (Ref. 8). LPN used only the default values given in Code 4.2, which of course can be overridden. The results are presented in Table 1.

Seven problems generated by the `lpgen` Code 4.1 were solved by the LPN Code 4.2. Problem size varied between 10^5 to 10^7 in the number of nonzero elements of the constraint matrix A . LPN solved all these problems in 11 to 26 iterations to an accuracy better than 10^{-13} . Comparing the times for the problems solved, CPLEX ran out of memory on two problems and was 2.8 to 34.2 times slower on the remaining five problems, with one problem (the third in Table 1), very poorly solved by CPLEX with a primal

Table 1. Comparison of LPN Algorithm 3.1 and CPLEX 6.5 (Ref. 8).

Problem Size \times Density $m \times n \times d$	LPN			Dual CPLEX 6.5	
	Time (sec)	Iter. #	Accuracy $\ x - z\ _\infty$	Time (sec)	Accuracy $\ x - r\ _\infty$
$(2.0 \times 10^6) \times 100 \times 0.05$	1041.8	26	1.1×10^{-14}	oom	oom
$(1.5 \times 10^6) \times 100 \times 0.05$	787.0	26	8.8×10^{-15}	oom	oom
$(1.0 \times 10^5) \times 1000 \times 0.1$	840.5	14	5.8×10^{-14}	28716.6	7.0
$(1.0 \times 10^5) \times 100 \times 1.0$	228.7	15	8.9×10^{-15}	905.0	7.5×10^{-14}
$(1.0 \times 10^5) \times 100 \times 0.1$	50.7	18	8.9×10^{-15}	238.4	5.3×10^{-12}
$(1.0 \times 10^4) \times 1000 \times 0.1$	417.5	11	5.1×10^{-14}	3513.2	5.4×10^{-12}
$(1.0 \times 10^4) \times 100 \times 0.1$	4.9	17	7.3×10^{-15}	13.8	1.3×10^{-11}

“oom” denotes out of memory. LPN time is total time: toc2 from Code 4.2.

solution error of 7.0. The dual CPLEX option was used throughout the numerical test problems because it gave the best results for the type of problems tested. The other options, primal and barrier, did worse. For example on the fifth test problem in Table 1 with dual CPLEX time of 238.4 seconds, the primal CPLEX took 2850.8 seconds, while the barrier CPLEX ran out of memory.

In order to show that LPN can handle also linear programs with non-unique solutions, as requested by one of the referees who noted also that the linear programs of Table 1 as generated by `lpgen` had unique feasible points as well, we perturbed the right-hand side b of (2) so that the number of active constraints at the solutions obtained by LPN and CPLEX were less than n . These results are presented in Table 2 and show that the accuracy of LPN has decreased. However, it is still capable of solving these problems to a reasonable accuracy and can still handle the two largest problem which CPLEX ran out of memory on.

Code 4.1. `lpgen` MATLAB lp Test Problem Generator.

```
%lpgen: Generate random solvable lp: min c'x s.t. Ax=<b;
      A: m-by-n
%Input: m,n,d(ensity); Output: A,b,c;
      (x,u): primal-dual solution
pl=inline('(abs(x)+x)/2'); %pl(us) function
tic; A=sprand(m,n,d); A=100*(A-0.5*spones(A));
u=sparse(10*pl(rand(m,1)-(m-3*n)/m));
x=10*spdiags((sign(pl(rand(n,1)-rand(n,1))))),0,n,n)
    *(rand(n,1)-rand(n,1));
c=-A'*u; b=A*x+spdiags((ones(m,1)-sign(pl(u))),0,m,m)
    *10*ones(m,1); toc0=toc;
format short e; [m n d toc0]
```

Table 2. Comparison of LPN Algorithm 3.1 and CPLEX 7.1 (Ref. 8) on linear programs with nonunique solutions.

Problem Size \times Density $m \times n \times d$	LPN				Dual CPLEX 7.1	
	Time (sec)	Iter. #	Accuracy $ c'z + b'v $ $\ (-Az + b)_+\ _\infty$	Act. Constr. #	Time (sec)	Act. Constr. #
$(2.0 \times 10^6) \times 100 \times 0.05$	3280.0	5	3.8×10^{-5} 4.1×10^{-6}	34	oom	oom
$(1.5 \times 10^6) \times 100 \times 0.05$	1899.1	5	1.3×10^{-4} 3.3×10^{-6}	33	oom	oom
$(1.0 \times 10^5) \times 1000 \times 0.10$	3014.5	7	5.7×10^{-5} 2.5×10^{-6}	240	15335.0	102
$(1.0 \times 10^5) \times 100 \times 1.00$	190.0	6	8.7×10^{-5} 2.5×10^{-6}	32	317.5	57
$(1.0 \times 10^5) \times 100 \times 0.10$	52.3	6	2.0×10^{-10} 7.3×10^{-7}	21	95.6	89
$(1.0 \times 10^4) \times 1000 \times 0.10$	48.2	8	1.4×10^{-7} 1.6×10^{-6}	43	2018.4	105
$(1.0 \times 10^4) \times 100 \times 0.10$	4.5	7	2.5×10^{-6} 1.8×10^{-6}	22	6.8	86

“oom” denotes out of memory.

Code 4.2. lpnewt1 MATLAB LPN Algorithm 3.1 without Armijo.

```

%lpnewt1: Solve primal LP: min c'x s.t. Ax=<b
%Via Newton for least 2-norm dual LP: max -b'v s.t. -A'v=c, v>=0
%Input: c, A, b, epsi, delta, tol, itmax;
      Output: v (12norm dual sol), z primal sol
epsi=1e-3; tol=1e-12; delta=1e-4;
      itmax=100; %default inputs
pl=inline('(abs(x)+x)/2'); %pl(us) function
tic; i=0; z=0; y=(A(1:n,:))'*A(1:n,:)+epsi*eye(n)\
      (A(1:n,:))'*b(1:n); %initial y
while (i<itmax & norm(y-z,inf)>tol & toc<1800)
      df=A'*pl((A*y-b))+epsi*c;
      d2f=A'*spdiags(sign(pl(A*y-b)),0,m,m)*A+delta*speye(n);
      z=y; y=y-d2f\df;
      i=i+1;
end
toc1=toc; v=pl(A*y-b)/epsi; t=find(v); z=A(t,:)\b(t); toc2=toc;
format short e;
      [epsi delta tol i-1 toc1 toc2 norm(x-y,inf) norm(x-z,inf)]
    
```

4.2. Machine Learning Test Problems. In this section, we test and compare LPN with CPLEX on four classification datasets from the UCI Machine Learning Repository (Ref. 22) and two publicly available datasets (Ref. 23). Again, we used LOCOP2.

The classification model that we shall employ here consists of a linear support vector machine (Refs. 24, 25, 27, 28) which we can state as the following linear program:

$$\min_{(w, \gamma, s, \epsilon) \in \mathbb{R}^{q+1+q+1}} e's + v\epsilon, \quad (40a)$$

$$\text{s.t.} \quad D(Bw - e\gamma) + e\epsilon \geq e, \quad -s \leq w \leq s, \quad \epsilon \geq 0, \quad r'w \geq 2. \quad (40b)$$

Here, e is a column vector of ones and v is a positive parameter that balances the model's ability to generalize to new unseen data (small v) versus empirical data fitting (large v). The matrices $B \in \mathbb{R}^{p \times q}$ and $D \in \mathbb{R}^{p \times p}$ constitute the problem data. The p rows of B represent p given data points in the input space \mathbb{R}^q , with each point belonging to class +1 if the corresponding diagonal element of the diagonal matrix D is +1, or class -1 if the corresponding diagonal element is -1. The vector r is the difference between the mean of the points in class +1 and the mean of the points in class -1. The linear program (40) generates a separating plane,

$$x'w = \gamma, \quad (41)$$

which separates approximately the points of classes +1 and -1. This separating plane lies midway between the parallel bounding planes

$$x'w = \gamma \pm 1. \quad (42)$$

With a maximum error of ϵ , the points of class +1 lie in closed halfspace

$$\{x | x'w \geq \gamma + 1\}. \quad (43)$$

Similarly, with maximum error of ϵ , the points of class -1 lie in the closed halfspace

$$\{x | x'w \leq \gamma - 1\}. \quad (44)$$

The error ϵ is minimized by the linear program (40) relative to $e's = \|w\|_1$. The last constraint of the linear program (40) excludes the degenerate solution $w = 0$ by implying the usually satisfied condition that the 1-norm distance $2/\|w\|_\infty$ (Ref. 29) between the two parallel bounding planes (42) does not exceed the 1-norm distance $\|r\|_1$ between the means of the two classes as follows:

$$\|w\|_\infty \cdot \|r\|_1 \geq r'w \geq 2. \quad (45)$$

Table 3. Comparison of LPN Algorithm 3.1 and CPLEX 6.5 (Ref. 8) on six publicly available machine learning problems.

Problem $m \times n, p \times q$	LPN		Dual CPLEX 6.5	
	Sec Infeas	Train Test	Sec Infeas	Train Test
Bupa Liver	0.11	68.12%	0.05	64.64%
$359 \times 14, 345 \times 6$	8.10e-3	66.71%	3.08e-19	63.92%
Pima Indians	0.21	76.69%	0.13	73.83%
$786 \times 18, 768 \times 8$	4.30e-3	76.20%	1.29e-18	73.39%
Cleveland Heart	0.14	87.54%	0.09	83.16%
$325 \times 28, 297 \times 13$	1.31e-2	84.74%	5.44e-18	81.69%
Housing	0.24	86.96%	0.13	83.60%
$534 \times 28, 506 \times 13$	9.80e-3	85.97%	1.02e-17	83.72%
Galaxy Bright	6.23	99.63%	0.79	99.15%
$2492 \times 30, 2462 \times 14$	8.03e-4	99.38%	5.37e-15	99.16%
Galaxy Dim	13.86	94.58%	1.57	90.96%
$4222 \times 30, 4192 \times 14$	1.50e-3	94.43%	1.81e-18	90.71%

$m \times n$ is the size of the matrix A of (2). $p \times q$ is the size of the matrix B of (40). “Infeas” denotes solution infeasibility (46). “Train” is training set correctness. “Test” is tenfold testing set correctness.

The linear program (40) was set up in the form of the linear program (2) with

$$n = 2q + 2 \quad \text{and} \quad m = 2p + 2q + 2.$$

The linear program (2) was solved by both LPN and CPLEX for six standard test problems varying in size between 297 points to 4192 points and for input spaces of dimensionality varying between 6 and 14. The results are given in Table 3. Because these are relatively small problems, CPLEX solved these problems in less time than LPN and with much greater accuracy as measured by the relative infeasibility of a solution point to (2),

$$\|(Ax - b)_+\|_1 / \|Ax\|_1. \tag{46}$$

However, correctness of the classifying separating plane on both the original data (training correctness) and on the tenfold testing (testing correctness) obtained by leaving one tenth of the data out for testing, repeating ten times and averaging, were higher for LPN than those for CPLEX on all six test problems. In addition, the LPN correctness values were comparable to those obtained by other methods (Refs. 2, 30, 31). These test correctness values are the key properties by which the machine learning classifiers are evaluated and are not always best when a very accurate solution to the linear program (40) is obtained. The explanation for this apparent paradox is that a highly accurate solution causes overfitting and may bias the classifier toward the

given empirical data at the expense of unseen data and hence results in poor generalization which translates into lower test set correctness as seen here for the CPLEX results. Application of Newton type methods have yielded some of the best test correctness results for classification problems (Refs. 2, 19, 32). The parameter value for ν for both LPN and CPLEX was the same and was chosen to optimize training set correctness. For all problems in Table 3, the value $\nu = 10^5$ was used. All parameters of LPN used were set at the default values given in the MATLAB Code 4.2.

5. Conclusions

We have presented LPN, a fast Newton algorithm for solving a class of linear programs. Computational testing on problems with $m \gg n$ and $m \geq n$ demonstrate the effectiveness of LPN in comparison with a state-of-the-art linear programming solver for this class of problems. Future work includes theoretical conditions under which LPN terminates without a stepsize in a finite number of steps, as it does in all the numerical examples presented here. Computational improvements might include novel ways of generating a primal solution from the dual least 2-norm solution, thus enabling LPN to handle all types of linear programs.

References

1. MANGASARIAN, O. L., *A Finite Newton Method for Classification Problems*, Report 01-11, Data Mining Institute, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 2001; see also *Optimization Methods and Software*, Vol. 17, pp. 913–929, 2002 and <ftp://ftp.cs.wisc.edu/pub/dmi/tech-reports/01-11.ps>.
2. FUNG, G., and MANGASARIAN, O. L., *Finite Newton Method for Lagrangian Support Vector Machine Classification*, Report 02-01, Data Mining Institute, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 2002; see also *Neurocomputing*, Vol. 55, pp. 39–55, 2003 and <ftp://ftp.cs.wisc.edu/pub/dmi/tech-reports/02-01.ps>.
3. MANGASARIAN, O. L., *Normal Solutions of Linear Programs*, *Mathematical Programming Study*, Vol. 22, pp. 206–216, 1984.
4. MANGASARIAN, O. L., and DE LEONE, R., *Error Bounds for Strongly Convex Programs and (Super) Linearly Convergent Iterative Schemes for the Least 2-Norm Solution of Linear Programs*, *Applied Mathematics and Optimization*, Vol. 17, pp. 1–14, 1988.
5. MANGASARIAN, O. L., and MEYER, R. R., *Nonlinear Perturbation of Linear Programs*, *SIAM Journal on Control and Optimization*, Vol. 17, pp. 745–752, 1979.

6. GOLIKOV, A. I., and EVTUSHENKO, Y. G., *Search for Normal Solutions in Linear Programming*, Computational Mathematics and Mathematical Physics, Vol. 14, pp. 1694–1714, 2000.
7. KANZOW, C., QI, H., and QI, L., *On the Minimum Norm Solution of Linear Programs*, Preprint, University of Hamburg, Hamburg, Germany, 2001; see also Journal of Optimization Theory and Applications (to appear) and <http://www.math.uni-hamburg.de/home/kanzow/paper.html>.
8. CPLEX Optimization, Incline Village, Nevada, *Using the CPLEX(TM) Linear Optimizer and CPLEX(TM) Mixed Integer Optimizer (Version 2.0)*, 1992.
9. CLARKSON, K. L., *Las Vegas Algorithms for Linear and Integer Programming*, Journal of the Association for Computing Machinery, Vol. 42, pp. 488–499, 1995.
10. KRISHNAN, K., and MITCHELL, J., *A Linear Programming Approach to Semidefinite Programming Problems*, Working Paper, Rensselaer Polytechnic Institute, Troy, NY, 2001.
11. PINAR, M. C., *Piecewise-Linear Pathways to Optimal Solution Set in Linear Programming*, Journal of Optimization Theory and Applications, Vol. 93, pp. 619–634, 1997.
12. HIRIART-URRUTY, J. B., STRODIOT, J. J., and NGUYEN, V. H., *Generalized Hessian Matrix and Second-Order Optimality Conditions for Problems with C^{L1} Data*, Applied Mathematics and Optimization, Vol. 11, pp. 43–56, 1984.
13. FACCHINEL, F., *Minimization of SC^1 Functions and the Maratos Effect*, Operations Research Letters, Vol. 17, pp. 131–137, 1995.
14. SMITH, P. W., and WOLKOWICZ, H., *A Nonlinear Equation for Linear Programming*, Mathematical Programming, Vol. 34, pp. 235–238, 1986.
15. FIACCO, A. V., and McCORMICK, G. P., *Nonlinear Programming: Sequential Unconstrained Minimization Techniques*, John Wiley and Sons, New York, NY, 1968.
16. BERTSEKAS, D. P., *Nonlinear Programming*, 2nd Edition, Athena Scientific, Belmont, Massachusetts, 1999.
17. MANGASARIAN, O. L., *Parallel Gradient Distribution in Unconstrained Optimization*, SIAM Journal on Control and Optimization, Vol. 33, pp. 1916–1925, 1995; see also <ftp://ftp.cs.wisc.edu/tech-reports/reports/1993/tr1145.ps>.
18. ARMIJO, L., *Minimization of Functions Having Lipschitz-Continuous First Partial Derivatives*, Pacific Journal of Mathematics, Vol. 16, pp. 1–3, 1966.
19. LEE, Y. J., and MANGASARIAN, O. L., *SSVM: A Smooth Support Vector Machine*, Computational Optimization and Applications, Vol. 20, pp. 5–22, 2001; see also Report 99-03, Data Mining Institute, University of Wisconsin, 1999 and <ftp://ftp.cs.wisc.edu/pub/dmi/tech-reports/99-03.ps>.
20. LUCIDI, S., *A New Result in the Theory and Computation of the Least-Norm Solution of a Linear Program*, Journal of Optimization Theory and Applications, Vol. 55, pp. 103–117, 1987.
21. MANGASARIAN, O. L., *Nonlinear Programming*, SIAM, Philadelphia, Pennsylvania, 1994.
22. MURPHY, P. M., and AHA, D. W., *UCI Machine Learning Repository*, 1992; see also www.ics.uci.edu/~mllearn/MLRepository.html.

23. ODEWAHN, S., STOCKWELL, E., PENNINGTON, R., HUMPHREYS, R., and ZUMACH, W., *Automated Star/Galaxy Discrimination with Neural Networks*, *Astronomical Journal*, Vol. 103, pp. 318–331, 1992.
24. MANGASARIAN, O. L., *Generalized Support Vector Machines*, *Advances in Large Margin Classifiers*, Edited by A. Smola, P. Bartlett, B. Schölkopf, and D. Schuurmans, MIT Press, Cambridge, Massachusetts, pp. 135–146, 2000; see also <ftp://ftp.cs.wisc.edu/math-prog/tech-reports/98-14.ps>.
25. CRISTIANINI, N., and SHAWE-TAYLOR, J., *An Introduction to Support Vector Machines*, Cambridge University Press, Cambridge, Massachusetts, 2000.
26. SMOLA, A., and SCHÖLKOPF, B., *Learning with Kernels*, MIT Press, Cambridge, Massachusetts, 2002.
27. VAPNIK, V. N., *The Nature of Statistical Learning Theory*, 2nd Edition, Springer, New York, NY, 2000.
28. BRADLEY, P. S., and MANGASARIAN, O. L., *Feature Selection via Concave Minimization and Support Vector Machines*, *Machine Learning, Proceedings of the 15th International Conference (ICML '98)*, Edited by J. Shavlik, and M. Kaufmann, San Francisco, California, pp. 82–90, 1998; see also <ftp://ftp.cs.wisc.edu/math-prog/tech-reports/98-03.ps>.
29. MANGASARIAN, O. L., *Arbitrary-Norm Separating Plane*, *Operations Research Letters*, Vol. 24, pp. 15–23, 1999; see also <ftp://ftp.cs.wisc.edu/math-prog/tech-reports/97-07r.ps>.
30. JOACHIMS, T., *Making Large-Scale Support Vector Machine Learning Practical*, *Advances in Kernel Methods – Support Vector Learning*, Edited by B. Schölkopf, C. J. C. Burges, and A. J. Smola, MIT Press, Cambridge, Massachusetts, pp. 169–184, 1999.
31. MANGASARIAN, O. L., and MUSICANT, D. R., *Lagrangian Support Vector Machines*, *Journal of Machine Learning Research*, Vol. 1, pp. 161–177, 2001; see also <ftp://ftp.cs.wisc.edu/pub/dmi/tech-reports/00-06.ps>.
32. FUNG, G., and MANGASARIAN, O. L., *A Feature Selection Newton Method for Support Vector Machine Classification*, Report 02-03, Data Mining Institute, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 2002; see also *Computational Optimization and Applications* (to appear) and <ftp://ftp.cs.wisc.edu/pub/dmi/tech-reports/02-03.ps>.