SIAM J. OPTIMIZATION
Vol. 2, No. 2, pp. 242–263, May 1992

004

A STRUCTURE-EXPLOITING ALGORITHM FOR NONLINEAR MINIMAX PROBLEMS*

ANDREW R. CONN† AND YUYING LI‡

Abstract. In this paper, some basic concepts are generalised which characterise the best linear Chebyshev approximation in one variable to general nonlinear minimax problems. A new method for solving a nonlinear minimax problem is presented, which exploits the structure and characterisation of the solution whenever possible. The algorithm is globally convergent with a superlinear convergence rate. Numerical results indicate the efficacy of the new method.

Key words. nonlinear Chebyshev approximation

AMS(MOS) subject classifications. 41A50, 65D99, 65F20, 65K05

1. Introduction.

We want to solve a *discrete nonlinear minimax problem*, which is written as

$$\min_{x\in\Re^n} \max_{i\in M} f_i(x), \tag{1.1}$$

where M is a finite index set. This is equivalent to finding the minimum value for the *maximum function* $\psi(x) = \max_{i\in M} f_i(x)$.

It is clear that a discrete Chebyshev problem

$$\min_{x\in\Re^n} \max_{1\le i\le m} |f_i(x)|, \tag{1.2}$$

which is a major class of discrete minimax problems, could be regarded as a special case of a general minimax problem (1.1) with

$$M = \{1, 2, \cdots, m, m+1, \cdots, 2m\}, \quad f_{i+m}(x) = -f_i(x), \quad i = 1, \cdots, m.$$

For simplicity, we describe our algorithm mainly in terms of the discrete Chebyshev problem (1.2) written in the form of (1.1). The extensions required for the general problem (1.1) are mentioned. In this paper, we are content to find a local minimum of (1.1) and we assume that a local minimum for (1.1) always exists. We also assume that each $f_i(x)$ is twice continuously differentiable.

Numerical methods for the discrete nonlinear Chebyshev/minimax problem are less prolific than for the linear problem. It is well known that the maximum function, $\psi(x) = \max_{i\in M} f_i(x)$, is not differentiable at kinks that arise whenever $f_i(x) = f_j(x)$, $i, j \in M$, $i \neq j$. Therefore, traditional gradient-type methods cannot be applied directly.

The existing methods are essentially based on successive linear programming or nonlinear programming techniques applied to an equivalent nonlinear programming problem. Examples include [1], [13], [20], [22], [23], [24], [27], [29], [30], [36], and [38].

* Received by the editors February 15, 1990; accepted for publication (in revised form) May 17, 1991.

† Thomas J. Watson Research Center, P.O.Box 218, Yorktown Heights, New York 10598. The research of this author was supported in part by Natural Sciences and Engineering Research Council of Canada grant A8639.

‡ Computer Science Department, Cornell University, Upson Hall, Ithaca, New York 14853. The research of this author was partially supported by the United States Army Research Office through the Mathematical Science Institute, Cornell University.

The classical Chebyshev theory provides us with characterisations for the best linear Chebyshev approximation. These properties uniquely determine a solution in many instances and thus requiring approximations with these special features is likely to result in a more efficient technique. Indeed, such has been the experience with classical Remez algorithms for best continuous/discrete linear Chebyshev approximation (see, for example, [35]).

In this paper, we first generalise the characterisation of the best linear Chebyshev approximation to a solution of a *nonlinear minimax problem.* The generalisation is useful computationally because we can force the approximate solutions to have these properties and thus expedite the solution-finding process. This is particularly beneficial for those problems arising from the discretisation of continuous approximation problems.

In developing our algorithm, we determine a suitable descent direction based on the structure of a solution (which consists of functions in general that are not necessarily close to the current maximum function). The new approach proposed is different from the existing methods in that one attempts to use the structure and characterisation of a solution of the minimax problem explicitly.

For clarity and brevity, we omit the proofs of some theorems. The interested reader is referred to [26] for details.

2. Structure of solutions to minimax problems. The continuous Chebyshev approximation problem on an interval $[a, b]$ can be described as

$$\min_{x\in\Re^n} \max_{t\in[a,b]} |f(t) - \phi(x,t)|,$$

where $f(t)$ and $\phi(x,t)$ are given functions.

Assume $\phi(x,t) = \sum_{i=1}^{n} x_i\phi_i(t)$. It is well known that, under the Haar condition, the absolute error function $|f(t) - \sum_{i=1}^{n} x_i\phi_i(t)|$ of the best linear Chebyshev approximation achieves the maximum value on $n+1$ points with the signs of the errors alternating [31]. Any ordered $n+1$ distinct points have been termed a *reference* and an approximation with the errors alternating signs on a reference has been called a *reference function* [35]. If a reference function has the same magnitude of errors on the reference, it is further called a *levelled* reference function.

The famous Remez algorithm finds the best Chebyshev approximation by constructing levelled reference functions at each step until a levelled reference function with the maximum error is obtained. For discrete linear Chebyshev approximations, the concept of reference and reference function has proven to be useful in developing computationally efficient algorithms (e.g., [4] and [6]).

Under some conditions, $|f(t)-\phi(x,t)|$ of the best nonlinear Chebyshev approximation achieves the maximum value at k points with the signs of the errors alternating [32]. Since the conditions are rather restrictive and k is not known a priori, there seems to be no computational algorithm that attempts to exploit the structure of the solution for a nonlinear Chebyshev problem.

In this section, we introduce the concepts of cadre and reference set for nonlinear minimax problems. They are generalisations of the corresponding concepts for linear Chebyshev problems.

DEFINITION 2.1. The vector set $\mathcal{C} = \{\nabla f_{i_j}\}_{j=0}^{l}$ is called a *cadre* if and only if:

1. $\operatorname{rank}([\nabla f_{i_0}, \cdots, \nabla f_{i_l}]) = l$;
2. for any $\{\nabla f_{j_1}, \cdots, \nabla f_{j_l}\} \subset \mathcal{C}$, $\operatorname{rank}([\nabla f_{j_1}, \cdots, \nabla f_{j_l}]) = l$.

This term was used by Descloux [18] to describe a linear Chebyshev solution when the Haar condition is not satisfied. A cadre can be equivalently defined by the following lemma.

LEMMA 2.2. $\mathcal{C} = \{\nabla f_{i_j}\}_{j=0}^{l}$ *is a cadre if and only if* $\operatorname{rank}(\mathcal{C}) = l$ *and there exist multipliers* $\{\lambda_i\}$ *such that*

$$\sum_{j=0}^{l} \lambda_j \nabla f_{i_j} = 0 \quad \textit{and} \quad \lambda_j \neq 0, \qquad j = 0, \cdots, l. \tag{2.1}$$

We refer to $\{\lambda_j\}$, *normalised by* $\sum_{j=0}^{l} \lambda_j = 1$, *if* $\sum_{j=0}^{l} \lambda_j \neq 0$ *and* $\lambda_0 = 1$ *otherwise, as* cadre multipliers. *The relation* (2.1) *is also called the* characteristic relation (*cf.* [31]).

Cadre multipliers are different from the Lagrangian multipliers used in optimization. The Lagrangian multipliers are usually associated with a stationary point and, under certain nondegeneracy assumptions, the nonzero multipliers are associated only with activities (see the following page for a definition of this term and the term ϵ-active). The cadre multipliers, however, are defined for any cadre and the functions in a cadre are not necessarily ϵ-active. Hence, we deliberately use the term cadre multipliers instead of just multipliers in order to differentiate them from the Lagrangian multipliers.

DEFINITION 2.3. The functions $\{f_{i_j}(x)\}_{j=0}^{l}$ are said to be locally forming a *reference set* of a minimax problem (1.1) if $\mathcal{C} = \{\nabla f_{i_j}\}_{j=0}^{l}$ is a cadre such that

1. The cadre multipliers $\{\lambda_j\}_{j=0}^{l}$ satisfy $\lambda_j > 0$, $j = 0, \cdots, l$;
2. The functions $\{f_{i_j}(x)\}_{j=0}^{l}$ all have the same sign.

The reference set is further called a *levelled reference set* if the value of each function is the same, viz.,

$$f_{i_j}(x) = f_{i_k}(x) \qquad \text{for any } i_j, i_k \in \mathcal{C}.$$

From the optimality conditions of (1.2) (e.g., [37]), we obtain an equivalent characterisation for a local minimum of (1.2) that relates to the structure of the best linear Chebyshev approximation.

THEOREM 2.4. *Suppose* x^* *is a local minimum for a minimax problem* (1.1). *Then, there exists a set of* $l+1$ *functions* $\{f_{i_j}(x)\}_{j=0}^{l}$, *which is a levelled reference set at* x^* *on the cadre* $\mathcal{C} = \{\nabla f_{i_j}(x^*)\}_{j=0}^{l}$ *with the maximum deviation.*

A reference set is a generalisation of the alternating sign property of a best Chebyshev approximation. Our experience with the numerical methods for linear l_∞ problems [6] suggests that it is very important to exploit computationally the above properties of a solution. The algorithm proposed in this paper is developed under this principle.

3. The model algorithm. The proposed algorithm is a descent method with a line search. The special features of the suggested algorithm, however, are that the search directions always decrease the maximum function and attempt to enforce the characterisation of a solution at the same time. Since a levelled reference set with the maximum deviation characterises a solution to a minimax problem, we attempt to compute the solution by constructing approximate solutions with such properties.

Assume $\mathcal{W} = \{i_0, i_1, \cdots, i_l\}$ is an index set and all the functions in $\mathcal{W}$ form a reference set that is not levelled. Denote

$$A = [\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_l}],$$
$$\Phi(x)^T = [f_{i_0}(x) - f_{i_1}(x), f_{i_0}(x) - f_{i_2}(x), \cdots f_{i_0}(x) - f_{i_l}(x)],$$

and $i_0 \in \mathcal{A}(x, 0)$. Here, $\mathcal{A}(x, 0)$ denotes the indices of the *active functions*, which are the functions achieving the maximum value at the current point x. In other words, $\mathcal{A}(x, 0) = \{i \in M | \psi(x) = f_i(x)\}$. More generally, we define the set of ϵ-active functions $\mathcal{A}(x, \epsilon)$ to be the set of functions that achieve the maximum deviation within a tolerance of ϵ, a small positive constant that may be reduced by the algorithm. That is, $\mathcal{A}(x, \epsilon) = \{i \in M | \psi(x) - f_i(x) \leq \epsilon\}$.

From the following two lemmas, it is possible to determine descent directions that attempt to construct a levelled reference set in the neighbourhood of a cadre or reference set.

LEMMA 3.1. *Suppose the functions in $\mathcal{W}$ form a reference set that includes all the current active functions. Then, the direction defined from $\mathcal{W}$ by*

$$v = -A(A^T A)^{-1}\Phi(x) \tag{3.1}$$

is a descent direction for all the active functions provided the reference set is not levelled.

In [29], a similar result, that the vertical direction v is a descent direction when the Lagrangian multipliers are nonnegative, is stated.

If a unit step along v is taken, $\Phi(x) + A^T v = 0$. Thus the functions in $\mathcal{W}$ would all have the same value as the representative function, a function chosen from $\mathcal{A}(x, \epsilon)$ at the start of the iteration, up to first order.

LEMMA 3.2. *Suppose $\mathcal{C} = \{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ is a nonreference set cadre with cadre multipliers $\{\lambda_j\}_{j=0}^{l}$ summing to one and $f_{i_0}(x)$ achieves the current maximum deviation for* (1.1). *Then, the direction v defined on $\mathcal{W} = \{i_0, i_1, \cdots, i_l\}$ by*

$$[\nabla f_{i_0} - \sigma_0\sigma_j \nabla f_{i_j}]^T v = -(f_{i_0} - \sigma_0\sigma_j f_{i_j}), \quad i_j \in \mathcal{W}, i_j \neq i_0, \ \sigma_j = \operatorname{sgn}(f_{i_j}), \tag{3.2}$$

decreases all the active functions, assuming $\mathcal{W}$ includes all the active functions at x.

If $\{f_{i_j}\}_0^l$ are linear functions at $x+v$, $\{f_{i_j}\}_0^l$ form a reference set. Thus, whenever the f_{i_j}'s do not constitute a reference set, moving along v, which is defined by (3.2),

$$[\nabla f_\mu - \sigma_0\sigma_j \nabla f_{i_j}]^T v = -(f_\mu - \sigma_0\sigma_j f_{i_j}), \qquad i_j \in \mathcal{W}, \quad i_j \neq i_0,$$

attempts to construct such a set.

We build up cadres using the concept of working sets. A working set is a function index set that includes all the indices of the current maximum functions. We emphasize, however, that the working set $\mathcal{W}$ is not generally an active set. In §5, we describe the details of setting up a working set.

The search direction is determined from the working set. If a cadre has not been located, in addition to decreasing the maximum function, the search direction is constructed to level the functions in the working set, when this is possible. The motivation behind this levelling comes from the fact that the structure of the solution requires the error curve to be levelled on the extreme points.

The suggested model algorithm is now outlined.

MODEL ALGORITHM

Step 1. Suppose an initial point x^0 is given. Set $k \leftarrow 0$.

Step 2. [Set up a working set]
The new working set $\mathcal{W}^k$ is determined. Check if there is a cadre $\mathcal{C}^k$ whose indices form a subset of $\mathcal{W}^k$. If there is no such cadre, go to Step 4.

Step 3. [Construct a levelled reference set]
Check reference set conditions. If the cadre corresponds to a reference set, compute a descent direction by levelling the reference set. Otherwise, find a descent direction that attempts to construct a reference set. Go to Step 5.

Step 4. [Descend and level]
A search direction d^k is found that decreases all the ϵ-active functions and levels the functions in the working set $\mathcal{W}^k$, if possible.

Step 5. [Line search]
A line search is performed on $\psi(x)$ along the direction d^k

$$x^{k+1} \leftarrow x^k + \lambda^k d^k; \quad k \leftarrow k+1.$$

Step 6. [Termination]
If optimal, stop. Otherwise, go to Step 2.

Step 3 of the model algorithm is one of the major parts in which the characterisation of the solution is exploited. From (3.1) and (3.2), we can compute a descent direction when a cadre is located (see also §6). Next, we discuss how to identify cadres (§4), how to construct a working set (§5), and how to compute a search direction when there is no cadre (§6). We also present details of the computation, including handling degeneracy (§8).

4. Identifying cadres. Given a set of functions $\{f_{i_0}, \cdots, f_{i_l}\}$, we discuss whether there exists a cadre within this set. We divide cadres into two types, depending upon whether

$$\sum_{j=0}^{l} \lambda_j = 1 \quad \text{or} \quad \sum_{j=0}^{l} \lambda_j = 0,$$

where $\{\lambda_j\}_{j=0}^{l}$ are cadre multipliers. The cadre that defines a reference set always belongs to the first type.

It is straightforward to prove the following lemma.

LEMMA 4.1. *Suppose* $\{\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_l}\}$ *are linearly independent. Then, the rank of the vector set* $\{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ *is at least* l.

The following lemma gives, under certain assumptions, necessary and sufficient conditions for the existence of a cadre with the sum of cadre multipliers being zero.

LEMMA 4.2. *Suppose* $A = [\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_{l-1}}]$ *is of full rank and that* $Z^T \nabla f_{i_0} \neq 0$, *where the columns of* Z *form a basis for the null space of* A^T. *Then, there exists a cadre* $\mathcal{C} \subseteq \{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ *with cadre multipliers summing to zero if and only if* $[\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_l}]$ *is rank deficient.*

Proof. Suppose $\mathcal{C} = \{\nabla f_{k_0}, \cdots, \nabla f_{k_\nu}\}$ is a cadre and $\{k_0, k_1, \cdots, k_\nu\} \subseteq \{i_0, i_1, \cdots, i_l\}$ with

$$\sum_{j=0}^{\nu} \lambda_j \nabla f_{k_j} = 0, \quad \sum_{j=0}^{\nu} \lambda_j = 0, \quad \lambda_j \neq 0, \quad j = 0, \cdots, \nu.$$

Then it is obvious that

$$\sum_{j=0}^{\nu} \lambda_j (\nabla f_{i_0} - \nabla f_{k_j}) = 0. \tag{4.1}$$

From (4.1) and the assumption that $\{\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_{l-1}}\}$ are linearly independent, we know that $i_l \in \{k_0, \cdots, k_\nu\}$. Hence, $\lambda_l \neq 0$ and we have

$$(\nabla f_{i_0} - \nabla f_{i_l}) = \sum_{j=1}^{l-1} \hat{\lambda}_j (\nabla f_{i_0} - \nabla f_{i_j}),$$

after padding with zeros if necessary. On the other hand, if we assume that $\{\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_{l-1}}\}$ are linearly independent and $\{\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_{l-1}}, \nabla f_{i_0} - \nabla f_{i_l}\}$ are linearly dependent, we have

$$\nabla f_{i_0} - \nabla f_{i_l} = \sum_{j=1}^{l-1} \hat{\lambda}_j (\nabla f_{i_0} - \nabla f_{i_j}). \tag{4.2}$$

From Lemma 4.1 and the assumption that A is full rank, we have that

$$\operatorname{rank}(\{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_{l-1}}\}) \geq l - 1.$$

Moreover, from $Z^T \nabla f_{i_0} \neq 0$, and the argument that follows, we can conclude that

$$\operatorname{rank}(\{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_{l-1}}\}) = l. \tag{4.3}$$

The above is true because, if $\{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_{l-1}}\}$ are linearly dependent, then there exist $\{\lambda_j\}$ that are not all zero such that

$$\sum_{j=0}^{l-1} \lambda_j \nabla f_{i_j} = 0.$$

If $\sum_{j=0}^{l-1} \lambda_j \neq 0$, without loss of generality, we can assume $\sum_{j=0}^{l-1} \lambda_j = 1$. Thus $\lambda_0 = 1 - \sum_{j=1}^{l-1} \lambda_j$. Hence

$$\nabla f_{i_0} = \sum_{j=1}^{l-1} \lambda_j (\nabla f_{i_0} - \nabla f_{i_j}).$$

We conclude that $Z^T \nabla f_{i_0} = 0$, which is a contradiction.

If $\sum_{j=0}^{l-1} \lambda_j = 0$, we have $\lambda_0 = -\sum_{j=1}^{l-1} \lambda_j$. Hence

$$\sum_{j=1}^{l-1} \lambda_j (\nabla f_{i_0} - \nabla f_{i_j}) = 0,$$

which is again a contradiction to the assumption that A is full rank.

Thus, using (4.2), we obtain

$$\sum_{j=0}^{l} \hat{\lambda}_j \nabla f_{i_j} = 0 \quad \text{and} \quad \sum_{j=0}^{l} \hat{\lambda}_j = 0, \tag{4.4}$$

where $\hat{\lambda}_0 = 1 - \sum_{j=1}^{l-1} \hat{\lambda}_j, \hat{\lambda}_l = -1$.

Define $\mathcal{C} = \{ \nabla f_{i_j} \mid \hat{\lambda}_j \neq 0 \ , \ j = 0, \cdots, \nu\}$. Using (4.4),

$$\text{rank}(\mathcal{C}) \leq |\mathcal{C}| - 1.$$

From (4.3), we know that

$$\text{rank}(\mathcal{C}) \geq |\mathcal{C}| - 1.$$

Hence

$$\text{rank}(\mathcal{C}) = |\mathcal{C}| - 1. \tag{4.5}$$

Moreover,

$$\hat{\lambda}_j \neq 0, \quad \nabla f_{i_j} \in \mathcal{C} \quad \text{with} \sum_{\nabla f_{i_j} \in C} \hat{\lambda}_j = 0.$$

Using Lemma 2.2, $\mathcal{C}$ is a cadre with the sum of the cadre multipliers being zero. □

Now, we present a lemma that tells us how to identify cadres with cadre multipliers summing to one.

LEMMA 4.3. *Suppose* $\{\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_l}\}$ *are linearly independent. Then there exists a cadre* $\mathcal{C} \subseteq \{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ *with cadre multipliers summing to one if and only if the orthogonal projected gradient,* $Z^T \nabla f_{i_0}$, *is zero, where*

$$A = [\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_l}], \qquad Z^T A = 0.$$

Proof. Since $\{\nabla f_{i_0} - \nabla f_{i_1}, \cdots, \nabla f_{i_0} - \nabla f_{i_l}\}$ are linearly independent, using Lemma 4.1,

$$\text{rank}(\{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}) \geq l. \tag{4.6}$$

The orthogonal projection of ∇f_{i_0} on the null space of A^T is $Z^T \nabla f_{i_0}$. The vector $Z^T \nabla f_{i_0}$ is zero if and only if there exist $\{\lambda_j\}_{j=0}^{l}$ such that

$$\lambda_0 \nabla f_{i_0} + \sum_{j=1}^{l} \lambda_j \nabla f_{i_j} = 0, \qquad \sum_{j=0}^{l} \lambda_j = 1. \tag{4.7}$$

Suppose (4.7) is satisfied. From (4.6) and (4.7), $\text{rank}(\{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}) = l$. Let $\mathcal{C} = \{\nabla f_{i_j} \mid \lambda_j \neq 0, j = 0, 1, \cdots, l\}$. Then, as in the argument for (4.5), $\mathcal{C}$ has rank $|\mathcal{C}| - 1$. From Lemma 2.2, $\mathcal{C}$ is a cadre. Moreover, the sum of the cadre multipliers is one.

On the other hand, if there is a cadre $\mathcal{C} \subseteq \{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ with cadre multipliers summing to one, then, following Lemma 2.2, there exist $\{\lambda_j\}$ such that (4.7) holds and then, $Z^T \nabla f_{i_0} = 0$. □

Lemmas 4.2 and 4.3 together enable us to determine whether there exists a cadre.

5. Establishment of the working set. A working set is a function index set, which is used to determine the current descent direction. Since we want the search direction to decrease all the ϵ-active functions, this working set $\mathcal{W}^k$ is chosen to include all the ϵ-active functions at the current point x^k. Nonetheless, there is flexibility in constructing such a set. We have chosen to build up the working set by selecting the

functions that determine the maximum function through several iterations. This is motivated by the fact that it is the extreme points that are important in determining the best approximation in the Chebyshev sense. Thus we require that

$$\mathcal{W}^k \subseteq \mathcal{W}^{k-1} \cup \mathcal{A}(x^k, \epsilon), \qquad \mathcal{W}^0 = \emptyset. \tag{5.1}$$

Moreover, the current ϵ-active functions are given priority over the functions in the old working set when forming the new working set.

However, since adjustment of the functions in the working set is necessary when the current working set is not approaching a reference set (essentially to account for the alternating sign property) we use $\hat{\mathcal{W}}^k$ to denote the set after possible modification and the rules for changing the set will be described precisely later. Thus, we more correctly require

$$\mathcal{W}^k \subseteq \hat{\mathcal{W}}^{k-1} \cup \mathcal{A}(x^k, \epsilon). \tag{5.2}$$

Assume, at the kth iteration, that a *representative function* $f_\mu(x)$, which can be any function $f_\mu(x)$ such that $\mu \in \mathcal{A}(x^k, \epsilon)$, is selected. Suppose $\mathcal{W}^k = \{\mu, i_1, \cdots, i_l\}$. The following Jacobian matrix corresponding to $\mathcal{W}^k$,

$$A^k = [\nabla f_\mu - \nabla f_{i_1}, \cdots, \nabla f_\mu - \nabla f_{i_l}], \tag{5.3}$$

is required *numerically* to have full rank. More specifically, our implementation accounts for this numerical rank. Conceptually it is equivalent to having some tolerance on the smallest singular value of A^k.

In implementation, we consider the projected gradient $Z^T \nabla f_\mu$ numerically zero if

$$\|Z^T \nabla f_\mu(x^k)\| \leq \tau_c^k,$$

where the columns of Z are an orthonormal basis for the null space of ${A^k}^T$ and τ_c^k is a small positive constant. Hence, if we identify cadres according to Lemma 4.3, we have a *near* cadre.

Since we need the QR decomposition (see, for example, [21, Chap. 6]) of the matrix A^k in computing the direction (see § 6), we build up the current working set $\mathcal{W}^k$ as follows.

CONSTRUCT $\mathcal{W}^k$:

Step 1. Set $Q \leftarrow I_{n\times n}$, $\mathcal{W}^k \leftarrow \{\mu\}$, where $\mu \in \mathcal{A}(x^k, \epsilon)$. $t \leftarrow 0$. $\hat{\mathcal{A}} = \mathcal{A}(x^k, \epsilon)$.

Step 2. If $\hat{\mathcal{A}} \setminus \mathcal{W}^k = \emptyset$, go to Step 3. Otherwise, let Q_2 be the last $n - t$ columns of Q and $j \in \hat{\mathcal{A}} \setminus \mathcal{W}^k$. If $\|{Q_2}^T(\nabla f_\mu - \nabla f_j)\| \leq \tau_0$, set $\hat{\mathcal{A}} = \hat{\mathcal{A}} \setminus \{j\}$, go to Step 2. Otherwise, go to Step 4.

Step 3. If $\hat{\mathcal{W}}^{k-1} \setminus \mathcal{W}^k = \emptyset$, stop. Otherwise, let Q_2 be the last $n-t$ columns of Q. If $\|Q_2^T \nabla f_\mu\| \leq \tau_c^k$, stop. Otherwise choose $j \in \hat{\mathcal{W}}^{k-1} \setminus \mathcal{W}^k$. If $\|{Q_2}^T(\nabla f_\mu - \nabla f_j)\| \leq \tau_0$, set $\hat{\mathcal{W}}^{k-1} = \hat{\mathcal{W}}^{k-1} \setminus \{j\}$ and go to Step 3. Otherwise, continue.

Step 4. Let $a = \nabla f_\mu - \nabla f_j$. Add the column a to A^k and update Q and R accordingly. Set:

$$A^k \leftarrow [A^k, a], \quad \mathcal{W}^k \leftarrow \mathcal{W}^k \cup \{j\}, \quad t \leftarrow t + 1.$$

Go to Step 2.

Thus the working set is the largest subset of $\hat{\mathcal{W}}^{k-1} \cup \mathcal{A}(x^k, \epsilon)$ (largest in the sense of the corresponding Jacobian matrix A^k being full rank), where the indices of the current ϵ-active functions have been entered preferentially.

Following the procedure of constructing a working set, it is clear that, if the current point is nondegenerate (a current point x^k is degenerate when there is a cadre $\mathcal{C} = \{\nabla f_{i_0}, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ such that $\{i_0, i_1, \cdots, i_l\} \subset \mathcal{A}(x^k, 0)$) and there is no cadre with cadre multipliers summing to zero, the Jacobian corresponding to all the ϵ-active functions is of full rank. Therefore

$$\mathcal{A}(x^k, \epsilon) \subseteq \mathcal{W}^k.$$

Moreover, if $\|Z^T \nabla f_\mu\| \leq \tau_c^k$, where $Z = Q_2$ for some Q, then a cadre (or a near cadre) with cadre multipliers summing to one is found.

6. Determining the search direction. Assume the working set at the current point x_c is

$$\mathcal{W}(x_c) = \{i_0, \cdots, i_l\} \quad \text{and} \quad \mu = i_0.$$

The desired search direction, in addition to being one of descent, attempts to enforce the characterisation of a solution.

Before a cadre with multipliers adding to one is located, we would like the search direction to decrease all the active functions and level all the functions in the working set, if possible. It is clear that $d = x - x_c$, where x attempts to solve

$$\begin{aligned} &\min_{x \in \Re^n} f_\mu(x) \\ &\text{subject to} \\ &f_\mu(x) - f_{i_j}(x) = 0, \qquad i_j \in \mathcal{W}(x_c), \end{aligned} \tag{6.1}$$

in the required direction. Note that μ is in fact a function of x_c and we use it to denote the current representative function as long as no confusion arises.

Dropping the subscript on x_c to simplify the description, one may approximate (6.1) as follows:

$$\begin{aligned} &\min_{d \in \Re^n} \nabla f_\mu(x)^T d + \tfrac{1}{2} d^T G d \\ &\text{subject to} \\ &\Phi(x) + A^T d = 0, \end{aligned} \tag{6.2}$$

where

$$\begin{aligned} A &= [\nabla f_\mu(x) - \nabla f_{i_1}(x), \nabla f_\mu(x) - \nabla f_{i_2}(x), \cdots, \nabla f_\mu(x) - \nabla f_{i_l}(x)], \\ \Phi(x) &= [f_\mu(x) - f_{i_1}(x), f_\mu(x) - f_{i_2}(x), \cdots, f_\mu(x) - f_{i_l}(x)]^T, \end{aligned}$$

and G is a matrix such that $Z^T G Z$ is positive definite, where the columns of Z form an orthonormal basis for the null space of A^T.

When close to a stationary point, $Z^T G Z$ is chosen to contain the curvature information of the functions in the working set in the null space of A^T (see §7 for details).

From the construction of the working set $\mathcal{W}(x)$, we know that A is of full rank. Following [11], the solution to (6.2) may be written as

$$\begin{aligned} d &= \hat{h} + v, \\ \hat{h} &= -Z(Z^T G Z)^{-1} Z^T (\nabla f_\mu(x) + Gv), \\ v &= -A(A^T A)^{-1} \Phi(x). \end{aligned}$$

It has been suggested in [11] that one could ignore the computation of $Z^T G v$ altogether without significantly effecting the rate of convergence. In this case, an approximate solution to (6.2) can be written as

$$d = h + v,$$

where

$$\begin{aligned} h &= -ZB^{-1} Z^T (\nabla f_\mu(x)), \\ v &= -A(A^T A)^{-1} \Phi(x), \end{aligned} \tag{6.3}$$

and

$$B = Z^T G Z.$$

It is clear that h is in the null space of A^T while v is in the range space of A. The direction in the null space of A^T will be called the *horizontal direction* and the direction in the range space of A will be called the *vertical direction.* We also point out that, given $\mathcal{W}$, Z, and B, the value of h and v is independent of the choice of μ (see [15] for details).

We now prove that a nonzero horizontal direction h is a descending direction for all the functions in $\mathcal{W}$.

LEMMA 6.1. *Assume $\mathcal{W}$ is the working set that defines the search direction. Assume further that B is positive definite and that there is no cadre $\mathcal{C} = \{\nabla f_{i_0}, \cdots, \nabla f_{i_l}\}$, with the cadre multipliers summing to one, such that $\{i_0, \cdots, i_l\} \subseteq \mathcal{W}$. Then the horizontal direction decreases all the functions in $\mathcal{W}$ equally (up to the first order); otherwise (i.e., there exists a cadre with the cadre multipliers summing to one), the horizontal direction h defined from $\mathcal{W}$ is zero.*

Proof. The horizontal direction defined in (6.3) is

$$h = -ZB^{-1} Z^T (\nabla f_\mu(x)), \qquad \mu = i_0,$$

where $Z^T Z = I_{n-l}$, $A^T Z = 0$. Since B is positive definite and

$$h^T \nabla f_\mu(x) = -(Z^T \nabla f_\mu(x))^T B^{-1} (Z^T \nabla f_\mu(x)),$$

it follows that

$$h^T \nabla f_\mu(x) < 0 \quad \text{iff } Z^T \nabla f_\mu \neq 0.$$

Since there is no cadre $\mathcal{C} = \{\nabla f_{i_0}, \cdots, \nabla f_{i_l}\}$ with the cadre multipliers summing to one such that $\{i_0, \cdots, i_l\} \subseteq \mathcal{W}$, we have, from the definition of $\mathcal{W}$ and Lemma 4.3, $Z^T \nabla f_\mu \neq 0$ and h is a descent direction for the representative function $f_\mu(x)$.

Furthermore, since

$$A^T h = 0 \quad \text{and} \quad \nabla f_{i_j}{}^T h = \nabla f_\mu{}^T h, \quad i_j \in \mathcal{W},$$

any function in the working set $\mathcal{W}$ will be decreased by the same amount (up to first order) as the representative function f_μ.

On the other hand, assuming there exists a cadre with cadre multipliers summing to one, by Lemma 4.3, the result follows. □

In conclusion, the horizontal direction h is a projection of the negative gradient of the representative function onto the null space of A^T. It is always a descent direction as long as $\mathcal{W}$ is not a cadre with cadre multipliers summing to one. As a descent direction, it decreases the functions in the working set by the same amount (up to first order). The horizontal direction h defined on the cadre with the cadre multipliers summing to one is always zero.

No cadre. When a cadre is not located, vertical directions are descent directions in most cases.

Whenever this is the situation, we perform the levelling process, i.e., set the search direction $d = v + h$. In the case in which the vertical direction is ascending, the vertical direction is discarded and the horizontal direction alone is taken as the search direction; specifically, we define

$$d^k = \begin{cases} h^k + v^k & \text{if } \nabla f_\mu{}^T v^k < 0, \\ h^k & \text{otherwise.} \end{cases} \tag{6.4}$$

Our numerical experience shows that an ascent vertical direction is a rare occurrence. This may be explained by the fact that the working set is constructed to approach a reference set. In the event that ascent does occur, we consider this as an indication that the working set is not approaching a reference set. This may be caused by some function, which will eventually not be maximum, being included in $\mathcal{W}^k$. Thus the next working set will not always include all the functions of the current working set; instead, we define

$$\begin{aligned} &\hat{\mathcal{W}}^k \leftarrow \mathcal{W}^k \setminus I^+, \quad \text{if } \nabla f_\mu^T v^k \geq 0, \text{ where} \\ &I^+ = \begin{cases} \{j_0\} & \text{if } \mathcal{A}(x^k, \epsilon) \subset \mathcal{W}^k \text{ and } f_\mu - f_{j_0} = \max_{j \in \mathcal{W}^k}(f_\mu - f_j); \\ \emptyset & \text{otherwise.} \end{cases} \end{aligned} \tag{6.5}$$

A cadre is located. If there exists a cadre with multipliers summing to zero, the cadre does not correspond to a reference set. In this case, although v corresponds to levelling, we emphasize decreasing the maximum function. In particular, it is not necessarily desirable to level functions that do not correspond to a reference set. Thus we simply take $d^k = h^k$. (Note that $h \neq 0$, since there is no cadre with cadre multipliers summing to one.)

If the functions in the working set, $\mathcal{W}^k$, form a (near) reference set, the vertical direction v^k defined by (6.3) attempts to level the functions in the working set while the horizontal direction h^k (again defined by (6.3)—$h^k = 0$ only if $\mathcal{W}^k$ contains an *exact* cadre with cadre multipliers summing to one) makes the gradients approach an *exact* cadre. From Lemma 3.1, v^k is a descent direction. Thus $d^k = h^k + v^k$ is a descent direction (note that h^k is a descent direction).

Suppose a cadre with multipliers summing to one has been located within the working set. Then the vertical direction v defined by (3.2) is a descent direction for the maximum function. Moreover, we can write (3.2) as

$$\begin{aligned} &\hat{A}v = -\hat{\Phi} \quad \text{where} \\ &\hat{A} = [\nabla f_\mu - \sigma_0\sigma_1 \nabla f_{i_1}, \cdots, \nabla f_\mu - \sigma_0\sigma_l \nabla f_{i_l}], \\ &\hat{\Phi} = [f_\mu - \sigma_0\sigma_1 f_{i_1}, \cdots, f_\mu - \sigma_0\sigma_l f_{i_l}]^T. \end{aligned} \tag{6.6}$$

We also modify the working set for the next iteration as follows. The cadre multipliers associated with the functions in the working set are used to construct the working set for the next iteration. The functions with positive multipliers are considered to be the functions which should be in the working set, i.e., the correct functions. For the functions with negative multipliers, we would like to put its negative function into the working set. However, because of nonlinearity and the fact that the cadre and reference set are both local properties, we prefer not to do so. Instead, the functions with negative multipliers are simply deleted from the working set, since the functions corresponding to negative multipliers will no longer remain ϵ-active when the direction v is taken and the multipliers sum to one. Thus we define

$$\hat{\mathcal{W}}^k \leftarrow \mathcal{W}^k \setminus \{\, i_j \mid \lambda_j < 0 \,\}. \tag{6.7}$$

The multipliers are thus used as a means to construct the working set and more than one function may be removed.

If the functions in the working set are all active and the multipliers sum to one, moving along the vertical direction initially decreases all the functions with the negative multipliers faster (up to first order) than those with positive multipliers. This comes from the following lemma (for the proof, see [14]).

LEMMA 6.2. *Suppose $\mathcal{W} = \{\mu, i_1, \cdots, i_l\}$ consists only of indices of the currently active functions. Assume further that $\mathcal{C} = \{\nabla f_\mu, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ is a cadre. Assume the direction v is determined from $\mathcal{W}$, as in* (3.2). *Then:*

1. *all the active functions with negative multipliers will be decreased more rapidly than all the other active functions, if the cadre multipliers sum to one, i.e., $\sum_{j=0}^{l} \lambda_j = 1$;*
2. *all the active functions are decreased equally (up to first order) provided the cadre multipliers sum to zero, i.e., $\sum_{j=0}^{l} \lambda_j = 0$.*

This corresponds to (possibly multiple) dropping of active functions for the equivalent nonlinear programming problem.

Now, consider a general nonlinear minimax problem written as

$$\min_{x \in \Re^n} \max_{i \in \{1, \cdots, m\}} f_i(x).$$

The search direction can be computed in exactly the same way except that the reference set, after a cadre has been located, could not be established as before. Since there exists no negative function of a given function, the vertical direction that determines which active functions should be dropped is not defined. Thus we now discuss how the definition of the vertical direction is modified for the general minimax problem.

If the current maximum deviation $\psi(x^k)$ is positive, we assume that for any given $f_i(x)$, there exists an imaginary $f_{i+m}(x) = -f_i(x)$. The working set $\mathcal{W}^k$ is chosen such that

$$-\psi(x^k) < f_{i_j}(x^k) \leq \psi(x^k) \quad \text{for any } i_j \in \mathcal{W}^k.$$

Hence locally we can treat the problem as a Chebyshev problem and the vertical direction, defined as for the Chebyshev problem, is a descent direction.

If the current maximum deviation $\psi(x)$ is nonpositive, we define a descent direction in a way similar to a general nonlinear programming approach [13]. In this case, if there exists some cadre multiplier that is negative, we simply remove the corresponding single function from the working set and update the projection matrix and recompute the search direction from the new projector. Under the assumption of linear independence, this will give a descent direction [13].

7. Approximation of the Hessian. In order to obtain a horizontal descent direction at each iteration B^k, an $(n-l)\times(n-l)$ matrix is assumed to be sufficiently positive definite.

For problems whose solutions are on a smooth valley, i.e., the number of active functions is less than $n+1$, the second-order information from the nonlinear active functions becomes significant for the fast final convergence of the algorithm. When close to x^*, B^k should be a good approximation to the projected Lagrangian Hessian ${Z^k}^T G^k Z^k$, where $G^k = \sum_{j=0}^{l} \lambda_j^k \nabla^2 f_{i_j}(x^k)$, the columns of Z^k form a basis for the null space of ${A^k}^T$, and ${\lambda_j}^k$ is an approximation to the Lagrangian multipliers (which are defined by the first-order optimality conditions of the equivalent nonlinear programming problem; see, for example, [13] or [38]).

If we assume the second-order sufficiency conditions hold at x^* and let λ^k be a good approximation to the cadre multipliers λ^* at *a solution* x^* (which are equal to the Lagrangian multipliers at a solution), then the matrix ${Z^k}^T G^k Z^k$, for x^k sufficiently close to x^*, is positive definite, as follows from continuity arguments.

A first-order method, for example, of [10], solves the problem whose solution is at a vertex (i.e., with $n+1$ linear independent activities) with a fast asymptotic rate of convergence since, once the correct activities are determined, one is merely using Newton's method (or a quasi-Newton method) to determine the unique intersection of these activities, with the corresponding quadratic (or superlinear) rate of convergence. First-order directions are usually good descent directions when one is far away from a stationary point and the computation of a first-order direction is cheaper than a second-order direction.

We choose to use the first-order direction if it gives a good improvement in the sense of constructing reference sets. Computationally, we consider that the first-order direction fails to improve the establishment of reference sets when the working set has not been changed for γ consecutive iterations (this may be a result of having the correct set but in this case it is reasonable to want to accelerate convergence by using a second-order direction). We arbitrarily set $\gamma = 3$ in our implementation. When failure occurs, we use the second-order information of the representative function or of all the functions in the working set, depending on how close we are to a stationary point of the subproblem.

Let *ibase* denote the number of consecutive iterations for which the working set remains unchanged. Suppose ρ is a small positive constant used to measure the closeness to a stationary point. The matrix G^k may be set up as follows:

$$G^k \begin{cases} \approx \nabla^2 f_\mu(x^k) & \text{if } ibase \geq \gamma \text{ and } \|{Z^k}^T \nabla f_\mu\| > \rho, \\ \approx \sum_{j=0}^{l} \lambda_j^k \nabla^2 f_{i_j}(x^k) & \text{if } ibase \geq \gamma \text{ and } \|{Z^k}^T \nabla f_\mu\| \leq \rho, \\ = I & \text{otherwise,} \end{cases} \tag{7.1}$$

where ${\lambda_j}^k$ is an approximation to the Lagrangian multipliers. We note that when $\|{Z^k}^T \nabla f_\mu\| \leq \rho$, it is reasonable to expect a suitable approximation to the Lagrangian multipliers.

Also, when $G^k = I$, the search direction is a first-order direction.

In our algorithm, however, we use a quasi-Newton method to update an approximation to the projected Hessian matrix B^k. Suppose Z^k is the orthogonal matrix such that ${Z^k}^T A^k = 0$, where A^k is defined as in (5.3). In the implementation, we have used the extended BFGS updating given below. B^k is initialised to be ${Z^k}^T G^k Z^k$

when necessary, where G^k is approximated according to (7.1) by finite differences. The extended BFGS updating follows:

$$B^{k+1} = B^k - \frac{1}{{s_r^k}^T B^k s_r^k} B^k s_r^k {s_r^k}^T B^k + \frac{1}{{y_r^k}^T s_r^k},$$

where

$$s_r^k = {Z^{k+1}}^T (x^{k+1} - x^k),$$
$$y_r^k = {Z^{k+1}}^T \nabla f_\mu(x^{k+1}) - {Z^k}^T \nabla f_\mu(x^k).$$

Assume B^k is positive definite. Then B^{k+1} remains positive definite if ${s_r^k}^T y_r^k > 0$. For unconstrained minimization, this condition is ensured by a line search. For constrained minimization, however, it cannot be satisfied in general. We have chosen to skip the update if the above condition is not satisfied.

8. Degeneracy. For a discrete Chebyshev problem, degeneracy handling is an important part of a useful algorithm. This is because, for example, in the linear case, it is not unusual for many residuals to achieve the maximum deviation. In this section, we discuss the handling of degeneracy in our algorithm.

We define a current point x^k to be degenerate when there is a cadre $\mathcal{C} = \{\nabla f_\mu, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ such that $\{\mu, i_1, \cdots, i_l\} \subset \mathcal{A}(x^k, 0)$.

Denote

$$\mathcal{W}^k = \{\mu, i_1, \cdots, i_l\}, \qquad A^k = [\nabla f_\mu - \nabla f_{i_1}, \cdots, \nabla f_\mu - \nabla f_{i_l}].$$

If x^k is a degenerate point, the following difficulty may occur. There is more than one cadre $\mathcal{C} = \{\nabla f_\mu, \nabla f_{i_1}, \cdots, \nabla f_{i_l}\}$ satisfying $\mathcal{W}^k \subset \mathcal{A}(x^k, 0)$. Thus it may not be possible to define a search direction such that it decreases the functions in all the cadres, although we know how to define a descending direction on one cadre.

If we consider the cadres that correspond to subsets of active functions, then there can be three types of degenerate points:

Type A. There only exist cadres with cadre multipliers summing to zero;

Type B. There exists a unique cadre and its cadre multipliers sum to one;

Type C. There exists more than one cadre and at least one with cadre multipliers summing to one.

A point x^* is a stationary point if and only if there exists at least one reference set consisting of active functions only.

We identify cadres by a tolerance of τ^k; the (numerical) degeneracy identified depends on the tightness of ${\tau_c}^k$. Thus when degeneracy is encountered, we reduce it by

$$\tau_c^{k+1} \leftarrow \frac{{\tau_c}^k}{2}. \tag{8.1}$$

Numerically, the degeneracy of Type A can only occur when $\|{Z^k}^T \nabla f_\mu^k\| > \tau_c^k$ and $\mathcal{W}^k \subset \mathcal{A}(x^k, \epsilon)$ [15]. For the degenerate points of Type A, there cannot be any reference set consisting of only the active functions. This is because, for any reference set, each of the corresponding cadre multipliers is positive and the sum of them is one. Thus the current point cannot be optimal. For this type of degeneracy, the horizontal direction h defined on the current working set decreases all the ϵ-active functions, up to first order, by the same amount.

Degeneracy of Type B occurs when $\|Z^{k^T}\nabla f_\mu^k\| \leq \tau_c^k$, $\mathcal{A}(x^k,\epsilon) = \mathcal{W}^k$, and there exists zero multipiers [15]. For the degenerate points of Type B, it is possible that a reference set exists within the active set. If there is such a reference set, then the current point is already a stationary point. Otherwise, since there exists a unique cadre, the vertical direction v defined on the cadre by (3.2) attempts to construct a levelled reference set. Moreover, other maximum functions not in the cadre can also be decreased at the same time.

If $\|Z^{k^T}\nabla f_\mu^k\| \leq \tau_c^k$ and $\mathcal{W}^k \subset \mathcal{A}(x^k,\epsilon)$, degeneracy of Type C occurs [15]. For the degenerate points of Type C, we do not know how to determine a descent direction without additional computation. Following a similar approach to [7] and [17], we solve the least squares problem:

$$\min_{\theta\in\Re^{l+1}} \left\| \sum_{j=0}^{l} \theta_j \nabla f_{i_j} \right\|_2$$

(8.2) subject to

$$\sum_{j=0}^{l} \theta_j = 1, \quad \theta_j \geq 0, \quad j = 0,\cdots,l, \quad \mu = i_0.$$

Assume $\{\lambda_j^k\}$ is the solution to (8.2). Analogous to the proof in [7], d^k defined by

$$d^k = -\sum_j \lambda_j^k \nabla f_{i_j} \tag{8.3}$$

is a descent direction unless $d^k = 0$, in which case we are optimal. Moreover, it is not difficult to prove that (8.2) can be solved via a least squares problem with only simple nonnegativity constraints [15].

9. Summary of the algorithm. Now we give a more detailed description of the algorithm.

Initialization: Suppose an initial point x^0 is given. Set $k \leftarrow 1$, $\epsilon \leftarrow \epsilon_0$, $\hat{\mathcal{W}}^0 \leftarrow \emptyset$.

Step 1. [QR decomposition]
Find the working set $\mathcal{W}^k \subseteq \hat{\mathcal{W}}^{k-1} \cup \mathcal{A}(x^k,\epsilon)$, Jacobian A^k, and its QR decomposition. Assume the columns of Z^k form a basis for the null space of A^{k^T}.
If $\mathcal{A}(x^k,\epsilon) \subseteq \mathcal{W}^k$ and $\|Z^{k^T}\nabla f_\mu\| \leq \tau_c^k$, go to Step 2;
If $\mathcal{A}(x^k,\epsilon) \subseteq \mathcal{W}^k$ and $\|Z^{k^T}\nabla f_\mu\| > \tau_c^k$, go to Step 3;
Set $\epsilon \leftarrow 0.1\epsilon$;
If $\mathcal{A}(x^k,\epsilon) \not\subseteq \mathcal{W}^k$ and $\|Z^{k^T}\nabla f_\mu\| > \tau_c^k$, go to Step 4;
If $\mathcal{A}(x^k,\epsilon) \not\subseteq \mathcal{W}^k$ and $\|Z^{k^T}\nabla f_\mu\| \leq \tau_c^k$, go to Step 5;
We note that the first and last instances imply that we have a cadre of type 1 ($\sum_{j=0}^{l} \lambda_j = 1$), and the third implies that we have a cadre of type 0 ($\sum_{j=0}^{l} \lambda_j = 0$).

Step 2. [Cadre "found" with $\sum_{i\in\mathcal{C}} \lambda_i = 1$]
If $\mathcal{W}^k$ is a reference set, obtain $B^k = Z^{k^T} G^k Z^k$, where G^k is defined as in (7.1); Compute the horizontal direction h^k and the vertical

direction v^k from (6.3); Set the search direction $d^k = h^k + v^k$ and $\hat{\mathcal{W}}^k \leftarrow \mathcal{W}^k$.

Otherwise, compute the vertical direction according to (6.6) and set $\hat{\mathcal{W}}^k$ using (6.7). Modify τ_c^k by (8.1) if degeneracy is encountered. Set $d^k = v^k$. Go to Step 6.

Step 3. [Cadre not found]

Obtain B^k as an approximation to ${Z^k}^T G^k Z^k$, where G^k is defined as in (7.1). Compute the horizontal direction h^k and the vertical direction v^k from (6.3). Compute the search direction d^k using (6.4). Set up $\hat{\mathcal{W}}^k$ according to (6.5). Go to Step 6.

Step 4. [Cadre "found" with $\sum_{i \in \mathcal{C}} \lambda_i = 0$]

Compute $d^k = -Z^k {Z^k}^T \nabla f_\mu^k$. $\hat{\mathcal{W}}^k \leftarrow \mathcal{W}^k$. Modify τ_c^k by (8.1) if degeneracy is encountered. Go to Step 6.

Step 5. [More than one cadre and at least one with $\sum_{i \in \mathcal{C}} \lambda_i = 1$]

Compute the search direction d^k using (8.3). Obtain $\hat{\mathcal{W}}^k$ from (6.7). Modify τ_c^k by (8.1).

Step 6. [Line search]

Perform a safeguarded line search. Set $k \leftarrow k + 1$. If $\|d^k\|_2 < \tau_s$ and $\mathcal{W}^k$ includes a levelled reference set, stop. Otherwise, go to Step 1. □

We use quotes around "found" to emphasize that τ_c^k is nonzero. The safeguards and details of the line search are given in [15].

10. Numerical testing. In this section, we compare the new algorithm with four other typical methods: [8], [13], [23], and [38].

The numerical results are for both minimax problems and discrete Chebyshev problems, all written in the form:

$$\min_{x \in \Re^n} \max_{i \in M} f_i(x). \tag{10.1}$$

The method of Conn. The method of [13] basically applies the active set strategy of nonlinear programming to the equivalent form of a minimax problem. It is a globally convergent algorithm with a superlinear convergence rate.

At each iteration, an equality-constrained quadratic programming subproblem is solved to determine the search direction. The subproblem is established upon all the current ϵ-active functions. The finite difference of the derivatives is used to approximate the second-order information.

This approach essentially corresponds to the sequential equality-constrained quadratic programming (EQP) approach for nonlinear programming problems, using projected Hessians. However, once the search direction is determined, the line search is done directly on the nondifferentiable maximum function $\psi(x)$.

Although there have been relatively fewer numerical results for general nonlinear minimax problems than for linear problems, to date, the available numerical results seem to indicate that the following method [23], which is a combination of a linear programming (LP) approach and a quasi-Newton method for a nonlinear system of equations, works well on most types of minimax problems.

The method of Hald and Madsen. At each iteration of the first stage, the method

of [23] requires an exact solution to a constrained linear minimax problem

$$\min_{d\in\Re^n} \max_{i\in[M]} \{f_i(x^k) + \nabla f_i(x^k)^T d\}$$
$$\text{subject to}$$
$$\|d\|_\infty \leq \Lambda^k$$

in order to find the search direction. A trust region method has been incorporated to ensure convergence.

If a solution is suspected of going through a smooth valley, i.e., the number of active functions at the solution is less than $n+1$, a switch to a second stage is made. Then a nonlinear system of equations established by the Kuhn–Tucker conditions for the active functions is solved by some quasi-Newton method.

The entire Lagrangian Hessian is approximated by some modified secant updates. It is possible for the maximum $\psi(x)$ to be increased. A return to the first stage might be necessary.

Under certain conditions, the method of [23] is globally convergent with a quadratic or superlinear final convergence, depending upon whether a Newton or a quasi-Newton method is involved.

The first stage of the method essentially corresponds to a sequential linear programming approach (SLP), stabilized via a trust region, for nonlinear programming problems.

The method of Womersley and Fletcher. The method of [38] is similar to that of [13]. It is a descent method which uses an active set strategy, a nonsmooth line search, and a quasi-Newton approximation to the projected Hessian of the Lagrangian function.

Global convergence of the algorithm has been proved. Under certain conditions, superlinear convergence occurs.

Like that of [13], this method could be considered as belonging to the class of sequential equality-constrained quadratic programming (EQP) approaches.

The method of Charalambous. In the approach of [8], the original minimax problem is defined as a modified least pth objective function which under certain conditions have the same optimum as the original problem.

10.1. Computational costs comparison. At each iteration, the methods of [13] and [38] and the new algorithm require the computation of a search direction obtained by solving an equality-constrained quadratic programming (EQP) or an equality-constrained linear programming (ELP). Comparable line searches have been used in the methods of Conn and Womersley and Charalambous and Fletcher, whereas Hald and Madsen used the trust region method. For our new algorithm, determining a cadre and dropping one function in the working set, when a nonreference set cadre is found, requires no extra work compared with the methods of [13] and [38]. When more than one function in the working set is dropped, an equivalent number of QR updates are required. Since these functions should be dropped and function evaluation typically is more expensive than a single QR update, in general, this extra work is well justified. The amount of computation per iteration required by the above three methods (i.e., [13] and [38] and the new method) is roughly the same.

The amount of work required by each iteration of [8] is roughly the same as performing a quasi-Newton step for an unconstrained function.

At each iteration of [23], in stage one, a linear programming problem of size at least $n \times |M|$ is solved up to optimality. At each iteration of stage two, if it is ever

entered, the computation required is similar to the methods of [13] and [38]. However, in general, most of the iterations are spent in stage one.

Loosely speaking, comparison of computational costs of one iteration of the new algorithm and that of [23] is similar to the comparison between one iteration of EQP and IQP methods.

A solution of EQP can be obtained by solving two linear systems of equations. The size of each linear system is at most n. A solution for IQP, however, usually requires iterative methods (i.e., inner iterations). Although the number of iterations are bounded by the number of unknowns and constraints, it is potentially very large and it could even become prohibitive for a discretised Chebyshev problem because the number of the constraints of its associated IQP can be much larger than those of the usual nonlinear programming problems.

Therefore, considering the amount of work required per iteration, the method of [23] is considerably more expensive than the others.

For nonlinear programming problems, the advantage of the IQP approach compared to EQP, however, has been the iterative search for the correct active set. Likewise, one would expect that the advantage of the method of [23] over that of [13] and [38] and the new algorithm is similar to that of the successive IQP method over the successive EQP approach for nonlinear programming problems; namely, it can identify the correct active set faster. This advantage probably is the case for the methods of [13] and [38]. The new algorithm, however, is not a pure active set method. It can also identify the correct active set quickly. It achieves this not by an iterative search but by recognising the structure of the optimum and constructively building up the reference set. Through exploiting the structure of the Chebyshev problem and minimax problem, we are able to retain the advantages of both the EQP approach and the IQP approach.

Finally, we remark that for a degenerate point of Type A or B, there is no extra work required compared with that for a nondegenerate point. For a degenerate point of Type C, we must solve a least squares problem with nonnegativity constraints.

10.2. Numerical results. We present some limited numerical results in this section.

For our numerical testing, the initial parameters required by the algorithm are set as

$$\tau_c^0 = 0.05, \quad \tau_0 = 10^{-12}, \quad \tau_s = \tfrac{1}{2}10^{-5}, \quad \rho = 0.5, \quad \epsilon_0 = 0.1.$$

The algorithm terminates when the following three conditions are satisfied:

1. $\|d^k\|_2 \leq \tau_s$;
2. $\mathcal{W}^k \subseteq \mathcal{A}(x^k, \epsilon)$;
3. $\lambda_j^k \geq 0$, for all $j \in \mathcal{W}^k$.

Thus, at termination, there exists, approximately, a levelled reference set with the maximum deviation.

The test problems include both nonlinear minimax problems and nonlinear Chebyshev problems.

We implicitly write a nonlinear Chebyshev problem

$$\min_{1 \leq i \leq m} |f_i(x)|$$

in the general minimax form

$$\min_{1 \leq i \leq 2m} f_i(x),$$

where $f_{i+m}(x) = -f_i(x)$, for $i = 1, \cdots, m$.

Consider the following nonlinear programming problem:

$$\begin{aligned} &\min_{x \in \Re^n} F(x) \\ &\text{subject to} \\ &g_i(x) \geq 0, \qquad i = 2, \cdots, m, \end{aligned}$$

and the minimax problem:

$$\begin{aligned} &\min_{x \in \Re^n} \max_{1 \leq i \leq m} f_i(x) \\ &\text{subject to} \\ &f_1(x) = F(x), \\ &f_i(x) = F(x) - \alpha_i g_i(x), \qquad 2 \leq i \leq m, \end{aligned}$$

where

$$\alpha_i \geq 0, \qquad 2 \leq i \leq m.$$

It is straightforward to show that for sufficiently large α_i, the optimum of the minimax problem coincides with that of the nonlinear programming problem (see [2]).

We have tested some nonlinear programming problems through the above transformation. The α parameter is set as

$$\alpha_i = 10.0, \qquad 2 \leq i \leq m,$$

which we know, a priori, is sufficiently large.

We have listed the results for the following minimax testing problems (their references are also indicated): Charalambous and Bandler 1 and Charalambous and Bandler 2 [9], Freudenstein and Roth [36], Colville problem 2 [12], Barrodale, Powell, and Roberts [5], Wong 1, Wong 2, and Wong 3 [10], Rosen and Suzuki [33], Rosenbrock [34], Transmission Problems [3], Davidon [16], Enzyme [25], El Attar [19], Hettich [36], Bard [36], Watson [36], and Osborne [36]. The starting points used are the same as that specified in the references.

The results for the problems Davidon, Enzyme, El Attar, and Hettich, under the column [23] are taken from [28], which describes essentially the same method as that of Hald and Madsen.

In Table 10.1, we report the number of function evaluations required by our new algorithm under the column NM. For each problem, we have used the nomenclature of the cited reference. The results of other methods, using a comparable stopping tolerance, are listed for comparison where available.

The column under the column *nact* indicates the number of maximum functions at the solution.

The Rosenbrock problem is degenerate at the solution. The Watson problem is degenerate at the starting point. The Watson problem with $n = 20$ is also degenerate at the solution obtained. For the other test problems, numerical degeneracy does not occur.

The reported results use extended BFGS updates. Similar results were obtained using exact derivatives. From the limited numerical results, we observe that, compared with [8], [13], and [38], the overall number of function evaluations required by the new

TABLE 10.1
Number of function evaluations: BFGS *updates*.

Problems	n	m	nact	NM	HM [23]	CN [13]	WF [38]	CL [8]
Charalambous & Bandler 1	2	3	2	11	11[a]	18	12	
Charalambous & Bandler 2	2	3	3	6	11[a]	8	6	
Freudenstein & Roth	2	2	2	11	15[a]			
Colville 2	15	21	12	49	41[a]	275	80	413
Barrodale, Powell et al.	5	21	5	21	10		38	
Wong1.1	7	5	3	25	23	106	53	107
Wong1.2	7	5	3	33	29	77	37	
Wong2	10	9	7	24	27[a]			120
Wong3	20	18	13	33	49[a]			318
Rosen & Suzuki	4	4	3	12	18	64	37	66
Rosenbrock	2	2	4	31	21			
Transmission 1	6	11	4	52	21	67		78
Transmission 2	6	11	4	25	46	80		
Davidon	4	20	3	20	27			
Enzyme	4	22	5	11	18			
El Attar	6	51	7	25	12			
Hettich	4	5	4	11	19[b]			
Bard	3	15	3	10	9[a]			
Madsen	2	3	2	17	13[a]			
Watson6	6	31	7	24	12[a]			
Watson20	20	31	39	22	39[a]			
Osborne	5	33	5	10	31[a]			

[a] The results are obtained by using the codes in [23].

[b] The algorithm stopped because of roundoff error without obtaining a solution.

algorithm is much less. We also recall that the amount of computation per iteration required by all but [8] to determine the search direction and stepsize are comparable. If one considers in more detail the number of function evaluations required and the size and complexity of the matrices being updated it would appear that the new method is more efficient than [8]. Hence, the new method appears to be more efficient than that of [8], [13], and [38].

The only method that seems to be competitive with the new algorithm is that of [23]. The number of function evaluations required by these two methods is comparable. However, we recall that the amount of remaining computation required per iteration demanded by the method of [23] is significantly more than the proposed method. Thus our new method still appears to be preferable.

We have also tested our new algorithm on a real application problem. The problem has 80 functions, in terms of a general minimax problem, with 40 variables. The number of activities at the solution is 39 (out of 80). Our algorithm solved it successfully in 50 function evaluations while the method of [23] failed to locate a solution.

11. Summary. The algorithm presented is a globally convergent algorithm with superlinear convergence rate [26]. It has been developed based on the principle that a minimax problem, in particular the Chebyshev problem, has special properties that can be computationally exploited in both the linear and nonlinear cases.

In this paper, we generalise the characterisation for a best linear Chebyshev approximation to nonlinear minimax problems. These generalisations are implementable computationally. We then present an algorithm which profits from this exploitation.

Typically, the algorithm attempts to find a cadre by focusing on the functions that have achieved maximum values through iterations, i.e., functions in working sets. These functions are then levelled by vertical directions whenever possible. If a reference set has been located, it is then levelled by vertical directions (which are descent directions) and thus a solution is quickly determined. If, however, the cadre does not correspond to a reference set, a descent direction is then defined as an attempt to construct one. Since, at any solution, there exists a *levelled* reference set with the *maximum* value, it is clear that the computational procedure is meaningful and we believe our numerical results indicate its promise.

We point out that it is possible for the Maratos effect to occur for the new algorithm as presently implemented. However, we have not experienced this effect during our numerical testing. Moreover, the algorithm can be slightly modified to guarantee that there is no Maratos effect. One only needs to reevaluate the functions at the point $x^k + h^k$ and compute the vertical direction using the updated values when one is close to a stationary point (see [15] for more details).

Finally, we point out that the algorithm can be extended to solve the constrained minimax problem (see [15] for more details).

REFERENCES

[1] D. H. Anderson and M. R. Osborne, *Discrete, nonlinear approximations in polyhedral norms: A Levenberg-like algorithm*, Numer. Math., 28 (1977), pp. 167–170.

[2] J. W. Bandler and C. Charalambous, *Nonlinear programming using minimax techniques*, J. Optim. Theory Appl., 13 (1974), pp. 607–619.

[3] J. W. Bandler and P. A. McDobald, *Optimization of microwave networks by razor research*, IEEE Trans. Microwave Theory Tech., 17 (1969), pp. 552–562.

[4] I. Barrodale and C. Phillips, *An improved algorithm for discrete Chebychev linear approximation*, in Proc. 4th Manitoba Conference on Numer. Math., University of Manitoba, Winnipeg, Canada, 1974, pp. 177–190.

[5] I. Barrodale, M. J. D. Powell, and F. D. K. Roberts, *The differential correction algorithm for rational l_∞-approximation*, SIAM J. Numer. Anal., 7 (1972), pp. 493–504.

[6] R. H. Bartels, A. R. Conn, and Y. Li, *Primal methods are better than dual methods for solving overdetermined linear systems in the l_∞ sense?*, SIAM J. Numer. Anal., 26 (1989), pp. 693–726.

[7] S. Busovača, *Handling degeneracy in a nonlinear l_1 algorithm*, Tech. Report CS-85-34, University of Waterloo, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 1985.

[8] C. Charalambous, *Acceleration of the least pth algorithm for minimax optimization*, Math. Programming, 17 (1979), pp. 270–297.

[9] C. Charalambous and J. W. Bandler, *Nonlinear minimax optimisation as a sequence of least pth optimization with finite values of p*, Internat. J. System Sci., 7 (1976), pp. 377–394.

[10] C. Charalambous and A. R. Conn, *An efficient method to solve the minimax problem directly*, SIAM J. Numer. Anal., 15 (1978), pp. 162–187.

[11] T. F. Coleman and A. R. Conn, *Nonlinear programming via an exact penalty function: Asymptotic analysis*, Math. Programming, 24 (1982), pp. 123–136.

[12] A. R. Colville, *A comparative study on nonlinear programming codes*, Tech. Report 320-2949, IBM Thomas J. Watson Research Center, Yorktown Heights, NY, 1968.

[13] A. R. Conn, *An efficient second order method to solve the (constrained) minimax problem*, Tech. Report CORR 79-5, Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Ontario, Canada, 1979.

[14] A. R. Conn and Y. Li, *Structure and characterization of discrete Chebyshev problems*, Tech. Report CS-88-39, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 1988.

[15] ———, *An approach to nonlinear l_∞ approximation*, in Proc. Fifth Mexican Workshop on Numerical Analysis, J. P. Hennart, ed., Society for Industrial and Applied Mathematics, Philadelphia, PA, 1991, pp. 346–365.

[16] W. C. DAVIDON, *A new least squares algorithm*, J. Optim. Theory Appl., 18 (1976), pp. 187–198.
[17] A. DAX, *A note on optimality conditions for the Euclidean multifacility location problem*, Math. Programming, 36 (1986), pp. 72–80.
[18] J. DESCLOUX, *Dégénéresence dans les approximations de Tschebysheff linéaries et discrètes*, Numer. Math., 3 (1961), pp. 180–187.
[19] EL-ATTAR, M. VIDYASAGAR, AND S. R. K. DUTTA, *An algorithm for l_1-approximation*, SIAM J. Numer. Anal., 16 (1979), pp. 70–86.
[20] R. FLETCHER, *A model algorithm for composite nondifferentiable optimization problems*, Math. Programming Stud., 17 (1982), pp. 67–76.
[21] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, MD, 1983.
[22] N. I. M. GOULD, *On solving three classes of nonlinear programming problems via simple differentiable penalty functions*, J. Optim. Theory Appl., 56 (1988), pp. 89–126.
[23] J. HALD AND K. MADSEN, *Combined* LP *and Quasi-Newton methods for minimax optimization*, Math. Programming, 20 (1981), pp. 49–62.
[24] S. P. HAN, *Variable metric methods for minimizing a class of nondifferentiable functions*, Math. Programming, 20 (1981), pp. 1–13.
[25] J. S. KOWALIK AND M. R. OSBORNE, *Methods for Unconstrained Optimization*, American Elsevier, New York, 1968.
[26] Y. LI, *An efficient algorithm for nonlinear minimax problems*, Ph.D. thesis, Computer Science Department, University of Waterloo, Waterloo, Ontario, Canada, 1988.
[27] K. MADSEN, *Minimax solution of nonlinear equations without calculating derivatives*, Math. Programming Stud., 3 (1975), pp. 110–126.
[28] ———, *Minimization of nonlinear approximation functions*, Ph.D. thesis, Technical University of Denmark, Lyngby, Denmark, 1985.
[29] W. MURRAY AND M. L. OVERTON, *A projected Lagrangian algorithm for nonlinear minimax optimization*, SIAM. J. Sci. Statist. Comput., 1 (1980), pp. 345–370.
[30] M. R. OSBORNE AND G. A. WATSON, *An algorithm for minimax approximation in the nonlinear case*, Comput. J., 12 (1968), pp. 63–68.
[31] M. J. D. POWELL, *Approximation Theory and Methods*, Cambridge University Press, Cambridge, U.K., 1981.
[32] J. R. RICE, *The Approximation of Functions* 2, *Nonlinear Theory*, Addison–Wesley, Reading, MA, 1969.
[33] J. B. ROSEN AND S. SUZUKI, *Construction of the nonlinear programming test problems*, Comm. ACM, 8 (1965), p. 113.
[34] H. H. ROSENBROCK, *An automatic method for finding the greatest or least value of a function*, Comput. J., 3 (1960), pp. 175–184.
[35] E. L. STIEFEL, *Numerical methods of Tchebycheff approximation*, in Proceedings of a Symposium, Mathematics Research Center, University of Wisconsin Press, Madison, WI, 1959.
[36] G. A. WATSON, *The minimax solution of an overdetermined system of non-linear equations*, J. Inst. Math. Appl., 23 (1979), pp. 167–180.
[37] R. S. WOMERSLEY, *Local properties of algorithms for minimizing nonsmooth composite functions*, Math. Programming, 32 (1985), pp. 69–89.
[38] R. S. WOMERSLEY AND R. FLETCHER, *An algorithm for composite nonsmooth optimization problems*, J. Optim. Theory Appl., 48 (1986), pp. 493–523.