

CONJUGATE GRIDS FOR UNCONSTRAINED
OPTIMISATION

D. Byatt, I. D. Coope and C. J. Price

*Department of Mathematics and Statistics
University of Canterbury
Private Bag 4800
Christchurch, New Zealand*

Report Number: UCDMS2002/10

August 2002

Keywords: Unconstrained optimisation, grid-based methods, pattern search, conjugate directions

CONJUGATE GRIDS FOR UNCONSTRAINED OPTIMISATION

D. BYATT, I. D. COOPE, AND C. J. PRICE

ABSTRACT. Several recent papers have proposed the use of grids for solving unconstrained optimisation problems. Grid-based methods typically generate a sequence of grid local minimisers which converge to stationary points under mild conditions.

The location and number of grid local minimisers is calculated for strictly convex quadratic functions in two dimensions with certain types of grids. These calculations show it is possible to construct a grid with an arbitrary number of grid local minimisers. The furthest of these can be an arbitrary distance from the quadratic's minimiser. These results have important implications for the design of practical grid-based algorithms.

Grids based on conjugate directions do not suffer from these problems. For such grids only the grid points closest (depending on the choice of metric) to the minimiser are grid local minimisers. Furthermore, conjugate grids are reasonably stable under mild perturbations so that in practice, only approximately conjugate grids are required.

1. INTRODUCTION

Several recent papers ([3, 4, 6, 12]) discuss the use of grids to solve the unconstrained optimisation problem

$$\min_{x \in \mathbb{R}^n} f(x)$$

where $f: \mathbb{R}^n \rightarrow \mathbb{R}$ and derivative information may not be explicitly available.

The basic idea behind grid-based or pattern search (as defined in [12]) methods is to find an analogue of a stationary point when the function is restricted to the nodes of a grid, for a sequence of progressively finer grids. It has been shown in [3, 12] that, under mild conditions, any limit point of such a sequence is a stationary point of the function. Both these papers include the case where convergence is proven for the sequence of grid points for which no adjacent grid point has lower function value. Such points are called *grid local minimisers* in [3] and *unsuccessful iterates* in [6, 12].

Although the number of grid local minimisers and their distance from the minimiser does not affect the theoretical properties of grid-based methods, it may have a significant effect in practice. If a sequence of grid local minimisers is located far from the minimiser the grid size may be reduced rapidly and prematurely. Under such conditions an algorithm may take steps which are tiny, requiring many iterations for any significant progress towards the minimiser [2] and convergence will be slow.

This research was financially supported by a Top Achiever Doctoral Scholarship.

For a given function and grid it is, in general, difficult to determine the number and position of grid local minimisers without evaluating the function at each of the grid points. For a strictly convex quadratic function one may intuitively think that a grid local minimiser is a grid point closest to the minimiser, regardless of the grid. This paper shows that even in two dimensions this idea is false for general grids, but true for some grids, including those based on conjugate directions.

Section 2 shows that only the grid points closest to the minimiser are grid local minimisers for grids based on conjugate directions. Further, the smallest angle between any pair of conjugate directions is calculated for strictly convex quadratic functions.

Section 3 gives an upper bound on the maximum distance from the minimiser to a grid local minimiser, and an explicit formula for the number of grid local minimisers for strictly convex quadratic functions in two dimensions, with a particular class of grid. The effect of rotation on the number of grid local minimisers for a particular strictly convex quadratic function and a selection of grids is also shown. How much a grid can be rotated without affecting a given grid local minimiser for strictly convex quadratic functions in two dimensions with a particular class of grid is also calculated.

Formal definitions of terms used throughout this paper are now presented.

1.1. Definitions. A *grid* $G_{\mathcal{V}}(h, x_0) \subset \mathbb{R}^n$ is defined by a set of n linearly independent basis vectors $\mathcal{V} = \{v_i\}_{i=1}^n$, a positive grid size parameter h and a point x_0 on the grid. The points of the grid are

$$G_{\mathcal{V}}(h, x_0) = \left\{ x \in \mathbb{R}^n : x = x_0 + h \sum_{i=1}^n \eta_i v_i, \eta_i \in \mathbb{Z} \right\}.$$

The grid size parameter h is adjusted from time to time in order to ensure that successive grids become finer in a manner needed to establish convergence [3]. The vectors v_i are referred to as the *grid directions*. A *grid line* is a line through one of the grid points in one of the grid directions. The vectors $\pm hv_i$ are the steps between *adjacent* grid points.

A grid point is a *grid local minimiser* for the function f if no adjacent grid point has lower function value. The term grid local minimiser will be abbreviated as *GLM*. If a point \check{x} on the grid $G_{\mathcal{V}}(h, x_0)$ is a GLM then

$$f(\check{x} \pm hv_i) \geq f(\check{x}) \quad \forall v_i \in \mathcal{V}.$$

The definition of a GLM is motivated by the fact that if \mathcal{V}^+ is a positive basis (see [5] for more details on the theory of positive bases) then [4] shows that

$$v_i^T \nabla f(x) \geq 0 \quad \forall v_i \in \mathcal{V}^+ \implies \nabla f(x) = 0.$$

The conditions which define a GLM are a finite difference approximation to this [3]. Although not relevant to the remainder of this paper, the grid definition above implicitly




Figure 1. A single cell of a grid showing that the closest grid point x_c to the minimiser x^* is not a grid local minimiser.

makes use of the *maximal* positive basis $\mathcal{V}^+ = \{\pm v_i\}$ with $2n$ elements (other positive bases are possible, see [3, 4, 5] for more details).

It is shown in [6] that at a GLM there is sufficient information to give an approximation to $\nabla f(x)$, or at least an upper bound on $\|\nabla f(x)\|$. See, for example, [4, 7, 9] for details on the practical problem of estimating derivative information using function values at grid points.

Given a strictly convex quadratic function with positive definite Hessian matrix B , a *B-conjugate grid* is one in which the grid directions are B -conjugate. Note that for a given quadratic function there are infinitely many B -conjugate grids. When the quadratic function is not in question, the shorter term *conjugate grid* may be used.

Depending on the shape of the grid, the closest (in the standard Euclidean metric) grid point to the minimiser x^* may not be a GLM. A two dimensional example is illustrated in Figure 1, which shows the contours of a strictly convex quadratic function and a single cell of a grid. The closest grid point to the minimiser is x_c and the GLM is \check{x} .

We now choose a new metric to measure the distance between two points relative to a grid $G_{\mathcal{V}}(h, x_0)$. Since the set $\mathcal{V} = \{v_i\}_{i=1}^n$ forms a basis for \mathbb{R}^n , any point $x \in \mathbb{R}^n$ can be represented uniquely as $x = h \sum_{i=1}^n \zeta_i v_i$, $\zeta_i \in \mathbb{R}$. The *grid norm* of such a point is defined as

$$\|x\|_{\mathcal{G}} = h \sum_{i=1}^n |\zeta_i| \|v_i\| \quad (1)$$

where $\|\cdot\|$ represents the standard Euclidean norm. The *grid distance* between any two points x, y is given by $\|x - y\|_{\mathcal{G}}$. This distance measures the shortest distance between x and y with travel restricted to the grid directions. Some authors refer to this as the *taxi-cab* metric. Note that there may not be a unique closest grid point to a given point $x \in \mathbb{R}^n$ due to the symmetry of the grid about x .




Figure 2. Rosenbrock's function and orthogonal grid showing the number and location of GLMs.

As the diagonals of each cell of a given grid are the same length using the grid distance metric, the *diameter* (of a cell) of the grid $G = G_{\mathcal{V}}(h, x_0)$ can be defined in a natural way as

$$\text{diam}(G) = h \sum_{i=1}^n \|v_i\|. \quad (2)$$

The following subsection shows that the number and location of GLMs depends on both the function and the choice of grid.

1.2. Position of grid local minimisers. Figure 2 shows the contours of Rosenbrock's function [11] and the number and position of GLMs with a square grid, centred on the origin with $h = 1/3$.

Figure 3 illustrates a strictly convex quadratic function and an orthogonal grid where the only GLM is the minimiser of the function. One may intuitively think that only the grid points closest to the minimiser will be GLMs for such nicely behaved functions. However, Figure 4 shows this is false, even in two dimensions. Figure 5 shows a two dimensional example of a conjugate grid.

All illustrations using a strictly convex quadratic function in two dimensions are based on the *example quadratic* function $q(x, y) = x^2 + 25y^2$. Occasionally the function used will be the example quadratic rotated about the origin (as in Figures 4 and 5).




Figure 3. Orthogonal grid with a strictly convex quadratic.

2. CONJUGATE GRIDS

For strictly convex quadratic functions, grids based on conjugate directions guarantee that only the closest grid points to the minimiser will be GLMs. Further, multiple GLMs exist only if there is symmetry of the grid about the minimiser. In this case, all GLMs will:

- Belong to the same cell of the grid.
- Have the same function value.
- Be the same (grid) distance from the minimiser.

Theorem 1. *For any quadratic function $q: \mathbb{R}^n \rightarrow \mathbb{R}$ with (symmetric) positive definite Hessian matrix B , minimiser x^* and a set of B -conjugate vectors $\mathcal{V} = \{v_i\}_{i=1}^n$ then for the grid $G = G_{\mathcal{V}}(h, x_0)$, if \check{x} is a GLM then*

$$\|x^* - \check{x}\|_G \leq \frac{1}{2} \text{diam}(G).$$

Proof. Write $x^* = \check{x} + h \sum_{i=1}^n \zeta_i v_i$, $\zeta_i \in \mathbb{R}$ and $q(x) = (x^* - x)^T B(x^* - x)/2 + c$ for some constant $c \in \mathbb{R}$. The grid point \check{x} is a GLM if and only if $q(\check{x}) \leq q(x) \forall x$ adjacent to \check{x} . The grid point x is adjacent to \check{x} if and only if $x = \check{x} \pm h v_k$ for some $k \in \{1, 2, \dots, n\}$. Hence, by conjugacy,

$$q(\check{x}) \leq q(x) \iff \zeta_k^2 v_k^T B v_k \leq (\zeta_k \pm 1)^2 v_k^T B v_k \iff |\zeta_k| \leq 1/2. \quad (3)$$




Figure 4. Orthogonal grid with a (rotated) strictly convex quadratic.

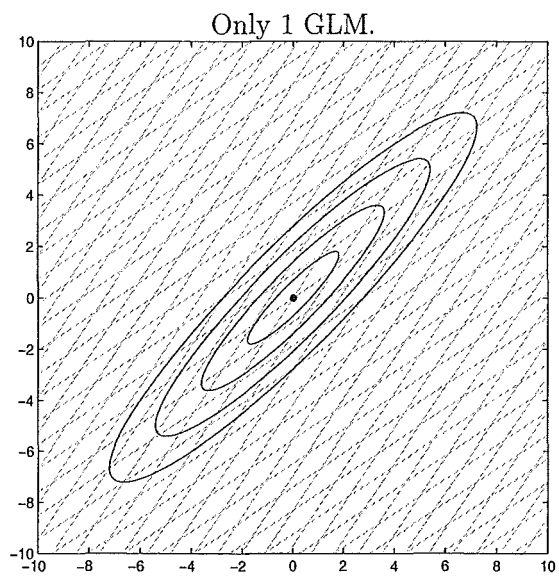


Figure 5. Conjugate grid with a (rotated) strictly convex quadratic.

Since equation (3) holds for every x adjacent to \check{x} , $|\zeta_i| \leq 1/2 \forall i \in \{1, 2, \dots, n\}$ and so by definitions (1) and (2)

$$\|x^* - \check{x}\|_G = h \sum_{i=1}^n |\zeta_i| \|v_i\| \leq \frac{1}{2} \text{diam}(G).$$

□

The following corollary follows immediately from Theorem 1.

Corollary 2. *If the conditions of Theorem 1 hold and \check{x} is a GLM then*

$$q(\check{x}) \leq q(y) \quad \forall y \in G$$

and

$$\|x^* - \check{x}\|_G \leq \|x^* - y\|_G \quad \forall y \in G.$$

Proof. Write $y = \check{x} + h \sum_{i=1}^n \eta_i v_i$, $\eta_i \in \mathbb{Z}$. Since $|\zeta_i| > |\zeta_i - \eta_i| \iff 0 < |\eta_i| < 1$ (which has no solutions for $\eta_i \in \mathbb{Z}$) it follows that $|\zeta_i| \leq |\zeta_i - \eta_i| \forall i \in \{1, 2, \dots, n\}$ and so

$$q(\check{x}) - q(y) = \frac{h^2}{2} \sum_{i=1}^n \zeta_i^2 v_i^\top B v_i - \frac{h^2}{2} \sum_{i=1}^n (\zeta_i - \eta_i)^2 v_i^\top B v_i \leq 0$$

and

$$\|x^* - \check{x}\|_G = h \sum_{i=1}^n |\zeta_i| \|v_i\| \leq h \sum_{i=1}^n |\zeta_i - \eta_i| \|v_i\| = \|x^* - y\|_G.$$

□

These results show that for a strictly convex quadratic function and a conjugate grid, every GLM is within half a cell diameter of the minimiser. Furthermore, if *any* GLM is found then no other grid point can be closer to the minimiser (using the grid distance metric), or have lower function value.

2.1. Smallest angle between conjugate directions. Although there are infinitely many sets of conjugate directions for a given strictly convex quadratic function, there exists a smallest (non-zero) angle between pairs of conjugate directions.

Theorem 3. *For any quadratic function $q: \mathbb{R}^n \rightarrow \mathbb{R}$ with (symmetric) positive definite Hessian matrix B , the smallest angle between any pair of B -conjugate directions is*

$$\theta_{\min} = 2 \tan^{-1}(\kappa^{-\frac{1}{2}}) \quad (4)$$

where $\kappa = \lambda_{\max}/\lambda_{\min}$ is the condition number of the matrix B whose largest eigenvalue is λ_{\max} and smallest eigenvalue is λ_{\min} .




Figure 6. Angle between conjugate directions.

Proof. Firstly the two dimensional case. Consider the ellipse $(x/a)^2 + (y/b)^2 = 1$ restricted to the first quadrant (symmetry takes care of the other cases) as shown in Figure 6. If $a = b$ the ellipse is a circle and the angle between any pair of conjugate directions is $\pi/2$. Suppose $a > b > 0$ and consider the line u through the origin and a point $P = (P_x, P_y)$ on the ellipse. If $P_x = 0$ or $P_x = a$ then the corresponding conjugate direction v is orthogonal to u . If $P_x \in (0, a)$ then u has slope $m_1 = \tan \theta_1 = P_y/P_x$ and the corresponding conjugate direction v has slope $m_2 = \tan \theta_2 = dy/dx|_P$. The angle $\theta = \theta_1 - \theta_2$ between u and v is given by

$$\tan \theta = \frac{m_1 - m_2}{1 + m_1 m_2}.$$

Minimising θ for $x \in (0, a)$ is equivalent to minimising $M = m_1 - m_2$ since $m_1 m_2 = -b^2/a^2$ is constant. Now M is minimised when $x = a/\sqrt{2}$ so that $m_1 = \tan \theta_1 = b/a$ and $m_2 = \tan \theta_2 = -b/a$.

Hence $\tan(\theta_{\min}/2) = b/a$. The geometric consequence of this is illustrated in Figure 7, which shows that minimum angle conjugate directions intersect the vertices of the bounding rectangle with sides parallel to the major and minor axes of the elliptical contours.

The ellipse $(x/a)^2 + (y/b)^2 = 1$ can be considered as a single contour of the two dimensional quadratic $q(x) = x^T B x/2$ for the (symmetric) positive definite matrix

$$B = \begin{pmatrix} 1/a^2 & 0 \\ 0 & 1/b^2 \end{pmatrix}$$

with eigenvalues $\lambda_{\min} = 1/a^2$, $\lambda_{\max} = 1/b^2$. Hence

$$\tan(\theta_{\min}/2) = \sqrt{\lambda_{\min}/\lambda_{\max}}.$$




Figure 7. Minimum angle conjugate directions.

Since the minimum angle between pairs of conjugate directions depends on the elongation of the elliptical contours, not on the orientation of the ellipses, and because any two conjugate directions define a plane, this result generalises nicely for higher dimensions. As the set of all possible smallest angles between pairs of conjugate directions is minimised when the (elliptical) contours are at their most elongated, result (4) holds for n dimensions. \square

Note that the conjugate grid shown in Figure 5 uses the smallest angle conjugate directions for the example quadratic function.

3. NUMBER OF GRID LOCAL MINIMISERS

In general, the number and position of GLMs is difficult to calculate. However for a strictly convex quadratic function in two dimensions, and a particular class of grid, it is possible to give an (attainable) upper bound for the maximum distance from a GLM to the minimiser (Theorem 4), and a formula for the total number of GLMs (Corollary 6).

All the grids considered in this section have two things in common:

- The minimiser of the quadratic function is a grid point.
- The principal axis of the quadratic function is parallel to one of the diagonals of the grid's cells (parallelograms in two dimensions).

Grids of this type will be referred to as *diagonally aligned grids*. Figure 8 illustrates a single cell of a general diagonally aligned grid.

For a diagonally aligned grid let θ_1, θ_2 be the internal angles of the cell between the grid directions and the principal axis of the quadratic function. Assume $\theta_1 \leq \theta_2$ so that, $\theta_1 \in (0, \pi/2)$ and $\theta_2 \in [\theta_1, \pi - \theta_1)$ (if $\theta_1 = 0$ the grid collapses to a line). Clearly $\|v_1\| \sin \theta_1 = \|v_2\| \sin \theta_2$ and the distance between grid points along the principal axis is $d = h(\|v_1\| \cos \theta_1 + \|v_2\| \cos \theta_2)$.

The only way that no component of a grid direction is towards the minimiser is if the grid line is orthogonal to the principal axis. As the grid directions must remain linearly independent, at least one of the grid directions cannot be orthogonal to the principal




Figure 8. Single cell of a general diagonally aligned grid.

axis. Hence, for the remainder of this discussion, assume $\theta_1 \leq \theta_2$ with $\theta_1 \in (0, \pi/2)$ and $\theta_2 \in [\theta_1, \pi/2]$.

Theorem 4. For any quadratic function $q: \mathbb{R}^2 \rightarrow \mathbb{R}$ with (symmetric) positive definite Hessian matrix B , and a diagonally aligned grid with parameters h , θ_1 , θ_2 , $\|v_1\|$, $\|v_2\|$, the maximum distance along the principal axis a GLM can be from the minimiser is

$$z_{\max} = \frac{h}{2} \|v_i\| \sin \theta_i (\kappa \tan \theta_i + \cot \theta_i), \quad i = \begin{cases} 1 & \text{if } \cot \theta_1 \cot \theta_2 \leq \kappa \\ 2 & \text{otherwise} \end{cases} \quad (5)$$

where $\kappa = \lambda_{\max}/\lambda_{\min}$ is the condition number of the matrix B whose largest eigenvalue is $\lambda_{\max} = 1/b^2$ and smallest eigenvalue is $\lambda_{\min} = 1/a^2$.

Proof. The contours of q are ellipses centred on x^* . With an appropriate co-ordinate system these contours have equation $(x/a)^2 + (y/b)^2 = (z/a)^2$ with $x \in [0, z]$ when restricted to the first quadrant (symmetry takes care of the other cases). Suppose $\check{z} = (z, 0)$ is a grid point on the principal axis and that C is the contour line of q which passes through \check{z} . If \check{z} is a GLM then none of the adjacent grid points has lower function value. That is, none of the grid points adjacent to \check{z} lie inside C .

The critical case which determines the maximum value for z (so that \check{z} is a GLM) occurs when none of the grid points adjacent to \check{z} lies inside C , but at least one of the adjacent grid points lie on C . A grid point will be referred to as *critical* when one of its adjacent grid points has the same function value. Suppose that for $i \in \{1, 2\}$, $z = z_i$ is a critical value of z so that \check{z} is a critical GLM. Furthermore suppose that $P = (P_x, P_y)$ is the corresponding critical grid point adjacent to \check{z} so that $P_x = z_i - h\|v_i\| \cos \theta_i$ and $P_y = h\|v_i\| \sin \theta_i$ (Figure 9 shows the case for $i = 1$). Then $a^2 P_y^2 = b^2(z_i^2 - P_x^2)$. Hence $z_i = h\|v_i\| \sin \theta_i (\kappa \tan \theta_i + \cot \theta_i)/2$. Therefore, as z increases from zero, the maximum distance between a GLM and the minimiser is determined by whichever grid point adjacent to \check{z} becomes critical first. Hence

$$z_{\max} = \min_{i \in \{1, 2\}} \{z_i\}.$$

It follows immediately that $z_1 \leq z_2 \iff \cot \theta_1 \cot \theta_2 \leq \kappa$ so that equation (5) holds. \square




Figure 9. Maximum distance grid local minimiser.

Perhaps even more important from a practical point of view is that there is also a minimum bound on the distance of the furthest GLM from the minimiser. This result is presented as the following corollary, which follows immediately from Theorem 4.

Corollary 5. *If the conditions of Theorem 4 hold then the distance of the furthest GLM from the minimiser is bounded below by*

$$\max(z_{\max} - d, 0) \quad (6)$$

where d is the distance between grid points along the principal axis of the quadratic function.

Proof. If $z_{\max} \geq d$, there must be a grid point on the principal axis whose distance from the minimiser is in the interval $(z_{\max} - d, z_{\max}]$. If $z_{\max} < d$, the only GLM is the minimiser. \square

The above results do not directly generalise for n dimensions as the plane with the most elongated contours may not contain *any* grid directions. However, Theorem 4 does give an upper bound for the n -dimensional case, which is attained when the “worst” grid directions coincide with the plane containing the most elongated contours.

Theorem 4 also shows that for any given strictly convex quadratic function, $z_{\max} \rightarrow \infty$ as $\theta_1 \rightarrow \pi/2$. Hence, it is possible to construct a diagonally aligned grid that has a GLM an arbitrary distance from the minimiser. This result has important practical implications. Several authors have proposed using the grid size parameter to determine a suitable stopping condition for grid-based algorithms. In fact, [12, p. 9] states:

... in the absence of any explicit higher-order information about the function to be minimized, it makes sense to terminate a generalized pattern search algorithm when Δ_k [the grid size parameter] is less than some reasonably small tolerance.

In fact, this is a common stopping condition for algorithms of this sort ...

The above result shows that unless some extra precautions are taken, using this condition alone does not guarantee that the minimiser (of even a strictly convex quadratic function) is within *any* given distance of the final iterate. A fact observed in [13, p. 196]:

For any non-derivative method, the issue of termination is problematical as well as highly sensitive to problem scaling. Since gradient information is unavailable, it is provably impossible to verify closeness to optimality simply by sampling f at a finite number of points.

The following example using a diagonally aligned grid illustrates the dangers of using only the grid size parameter as a stopping criterion.

Example 1. If $V = [v_1, v_2]$ is the matrix whose columns are the grid directions then a practical (convergent) grid-based algorithm will require $\det(V) \geq \delta$ for some (generally quite small) $\delta > 0$. Suppose the grid directions are $v_1 = [\delta, 1]^T$ and $v_2 = [0, 1]^T$ so that $\|v_1\| = 1/\sin \theta_1$, $\|v_2\| = 1$ and $\cot \theta_1 = \delta$. Since $\det(V) = \delta$, the grid is acceptable, and because $\cot \theta_1 \cot \theta_2 = 0$,

$$\begin{aligned} z_{\max} &= \frac{h}{2} \|v_1\| \sin \theta_1 (\kappa \tan \theta_1 + \cot \theta_1) \\ &= \frac{h}{2\delta} (\kappa + \delta^2) \\ &\approx \frac{h\kappa}{2\delta} \quad (\text{for small } \delta). \end{aligned}$$

If $\kappa = 100$ and $\delta = 10^{-6}$ then $z_{\max} \approx 5 \times 10^7 h$. So that even for a strictly convex quadratic function in just two dimensions, with condition number 100, an algorithm with stopping criterion based solely on h , could terminate when the distance from the minimiser to the current iterate is over seven orders of magnitude larger than h . Note that $\delta = 10^{-6}$ is not an overly demanding choice, and could be much smaller in practice. The convergent variant of the Nelder-Mead algorithm described in [10], for example, uses 10^{-18} as the lower bound on the linear independence of the simplex directions.

Clearly this is a bad grid for this problem. However, such information is rarely available in advance (compare the grids used in Figures 4 and 5, for example). Furthermore, whilst the transformation $x' = V^{-1}x$ orthogonalises the grid, the condition number of the transformed problem increases so much, that $z'_{\max} \approx 2 \times 10^{14} h$.

Forcing the grid to be orthogonal does not produce good results in practice (consider the performance of the method of alternating variables (which no one seems to want to take responsibility for), or the algorithm proposed by Hooke and Jeeves [8]). The above results go some way to explaining why. With a regular orthogonal grid $\tan \theta_1 = 1 = \tan \theta_2$, and so $z_{\max} \approx h\kappa/2$. Hence an algorithm could terminate when the distance from the minimiser to the current iterate is orders of magnitude larger than h .

These results suggest the need for stopping criteria of grid-based algorithms to be based on some measure of the conjugacy of the grid directions in addition to the grid size parameter.

Theorem 4 also provides the framework for the following corollary.

Corollary 6. For any strictly convex quadratic function $q: \mathbb{R}^2 \rightarrow \mathbb{R}$, and diagonally aligned grid with parameters $h, \theta_1, \theta_2, \|v_1\|, \|v_2\|$ the total number of GLMs is

$$1 + 2 \left\lfloor \frac{z_{\max}}{d} \right\rfloor \quad (7)$$

where $d = h(\|v_1\| \cos \theta_1 + \|v_2\| \cos \theta_2)$ is the distance between grid points along the principal axis.

Proof. Theorem 4 shows that if $\check{z} = (z, 0)$ is a grid point on the principal axis with $|z| \leq z_{\max}$ then \check{z} is a GLM as all adjacent grid points lie either outside, or on, the contour line through \check{z} . However, if $|z| > z_{\max}$ then at least one grid point adjacent to \check{z} will lie inside the contour line though \check{z} . Hence for a grid point \check{z} on the principal axis, \check{z} is a GLM $\iff |z| \leq z_{\max}$. Since the minimiser is a grid point, and by symmetry, the total number of GLMs is given by equation (7). \square

This result shows that:

- It is possible to construct a diagonally aligned grid with an arbitrary number of GLMs for any strictly convex quadratic function.
- The number of GLMs is independent of the grid size parameter.

The results of Corollary 6 do not directly generalise to higher dimensions. In n dimensions, GLMs may lie in some bounded region of an $n - 1$ dimensional hyperplane. In this situation the parameter z_{\max} corresponds to the size of the bounded region. The following example illustrates this situation in three dimensions.

Example 2. Consider the function $q(x, y, z) = x^2 + y^2 + z^2 + 2M(x + y + z)^2$, for some large positive integer M , and a (unit) cubic grid centred on the origin and aligned with the co-ordinate axes. If $\check{x} = (x_1, y_1, z_1)$ is a GLM on the plane $P: x + y + z = 0$ and x is a grid point adjacent to \check{x} then $x = (x_1 \pm 1, y_1, z_1)$ or $x = (x_1, y_1 \pm 1, z_1)$ or $x = (x_1, y_1, z_1 \pm 1)$. Suppose $x = (x_1 \pm 1, y_1, z_1)$, then $q(\check{x}) \leq q(x) \iff |x_1| \leq M$ (by symmetry, the same bound also applies to y_1 and z_1). Hence, grid points that lie in this bounded region of the plane P are GLMs. For this example, there are approximately $3M^2$ GLMs on the plane P , the furthest of which is a distance $M\sqrt{2}$ from the minimiser. Note that there may be other GLMs that do not lie on the plane P .

Clearly there can be very many more GLMs in higher dimensions. To make matters even worse, as the dimensionality increases, a greater proportion of these GLMs will be close to the maximum distance from the minimiser—a very bad situation in practice.

3.1. Effect of rotation on the number of GLMs. Figures 11-14 show the effects of rotation on the number of GLMs for the example quadratic function and each of the four sample grids described below. The rotations are about the minimiser in 0.1° increments. Note that since the contour lines for the example quadratic function are symmetric about

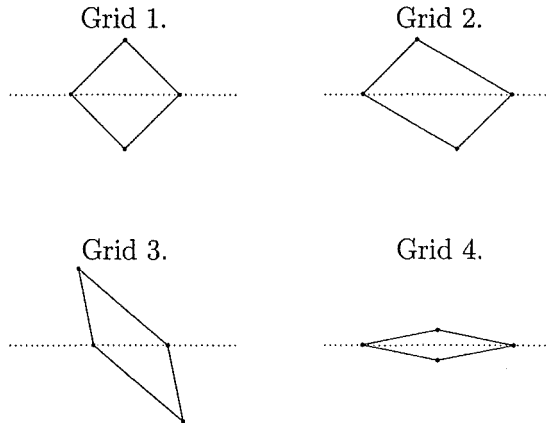


Figure 10. Single cells for each of the four sample grids.

the principal axis, the number of GLMs will be periodic, with period at most π radians. The period will be less if the grid is symmetric about a rotation of less than π (e.g. $\pi/2$) radians. Figure 10 shows a single cell for each of the four sample grids. For each of these grids, $h = 1$ and $\|v_2\| = 1$, so that $\|v_1\| = \sin \theta_2 / \sin \theta_1$. The remaining grid parameters are:

- Grid 1 (Regular orthogonal grid): $\theta_1 = \pi/4 = \theta_2$.
- Grid 2 (Miscellaneous grid): $\theta_1 = \pi/6$, $\theta_2 = \pi/4$.
- Grid 3 (Random grid): $\theta_1 = 0.6985$, $\theta_2 = 1.7722$.
- Grid 4 (Minimum angle conjugate grid): $\theta_1 = \tan^{-1}(b/a) = \theta_2$.

Figures 11-14 show that even for a function as simple as the example quadratic, a relatively small rotation can dramatically alter the number of GLMs. This highlights the difficulty in finding a general formula for the number of GLMs, as such a formula would have to reproduce this behaviour.

For each of the Figures 11-14 there is a relatively large angle of between about one half and one radian for which there is only one GLM. Although the size of this region of stability is quite large in two dimensions, the relative size of this n -dimensional cone of stability will, in general, tend to zero (rapidly) as the dimensionality increases. Under such circumstances, the probability that a general grid is aligned in such a way to produce few GLMs also tends to zero as the dimension of the problem increases—unless the grid is based on conjugate directions. In this case, Theorem 1 guarantees the existence of only one (up to symmetry) GLM. Note however, Figure 14 shows that if a conjugate grid becomes misaligned, the number of GLMs can increase dramatically. In fact, rotation of

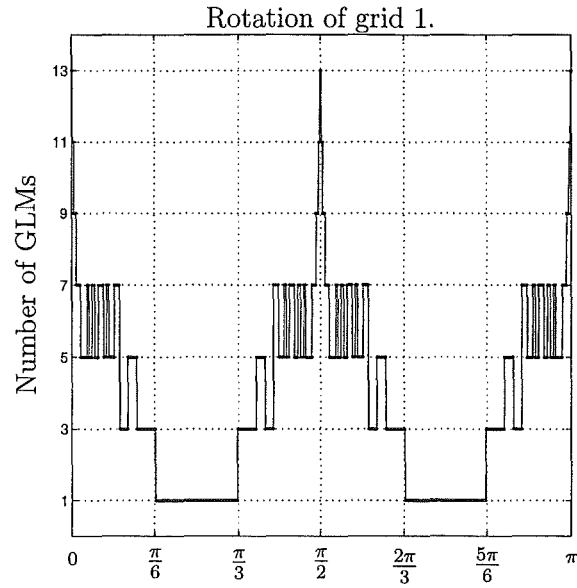


Figure 11. Rotation of grid 1 about the minimiser.

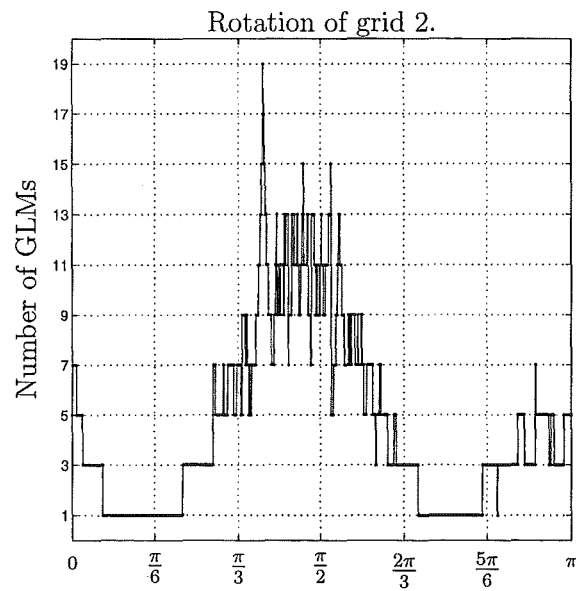


Figure 12. Rotation of grid 2 about the minimiser.

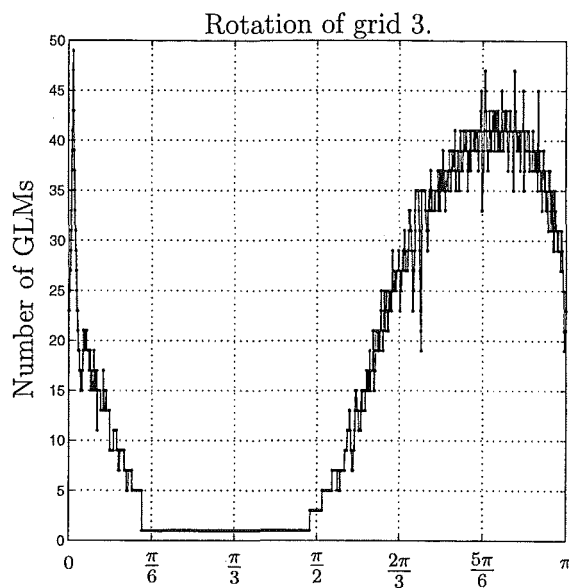


Figure 13. Rotation of grid 3 about the minimiser.

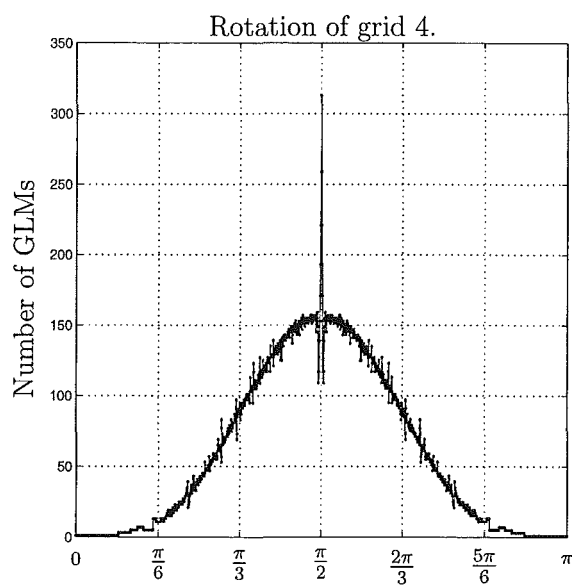


Figure 14. Rotation of grid 4 about the minimiser.

the minimum angle conjugate grid produced far more GLMs than rotation of any of the other sample grids.

3.2. Maximum rotation without affecting a GLM. This section examines how stable a given GLM is under rotation (about the minimiser) for a given diagonally aligned grid and a strictly convex quadratic function in two dimensions. Since any diagonally aligned grid with only one GLM (the minimiser), and functions with circular contours, are unaffected by *any* rotation (about the minimiser) they will be excluded from the following discussion. The following theorem gives a formula for the largest rotation of the grid, relative to the function, before a given GLM becomes critical.

Theorem 7. *For any quadratic function $q: \mathbb{R}^2 \rightarrow \mathbb{R}$ with (symmetric) positive definite Hessian matrix B with smallest eigenvalue $\lambda_{\min} = 1/a^2$ and largest eigenvalue $\lambda_{\max} = 1/b^2$, with $a > b > 0$, and diagonally aligned grid with parameters h , θ_1 , θ_2 , $\|v_1\|$, $\|v_2\|$, the maximum rotation ω_{\max} , of the grid, before a GLM $P_0 = (z_0, 0)$ on the principal axis becomes critical is given by*

$$\omega_{\max} = \min_{i \in \{1,2\}} \{\omega_i\} \quad (8)$$

where ω_i is the solution to

$$z_0^2 \sin^2 \omega_i - r_i^2 \sin^2 \psi_i = b^2 \left(\frac{r_i^2 - z_0^2}{a^2 - b^2} \right)$$

for $z_{\min} \leq z_0 \leq z_{\max}$, and where

- (a) $\psi_i = \sin^{-1}(h\|v_i\| \sin \theta_i / r_i) - \omega_i$,
- (b) $r_i^2 = z_0^2 + h^2\|v_i\|^2 - 2h\|v_i\|z_0 \cos \theta_i$,
- (c) $z_{\min} = \min_{i \in \{1,2\}} \{h\|v_i\| \cos \theta_i\}$.

Proof. The contours of q are ellipses centred on x^* . With an appropriate co-ordinate system these contours have equation $(x/a)^2 + (y/b)^2 = c^2$ with $x \in [0, ac]$ when restricted to the right half-plane (again, symmetry takes care of the other cases). Suppose C_0 is the contour of q which passes through the GLM $P_0 = (z_0, 0)$ on the principal axis of q (see Figure 15). If P_0 is rotated by ω_i to point $P_1 = (z_0 \cos \omega_i, -z_0 \sin \omega_i)$, then the contour C_1 of q passing through P_1 is given by $a^2 z_0^2 \sin^2 \omega_i = b^2 (z_1^2 - z_0^2 \cos^2 \omega_i)$ so that

$$z_1^2 = z_0^2 \left(\frac{a^2}{b^2} \sin^2 \omega_i + \cos^2 \omega_i \right).$$

For each $i \in \{1, 2\}$, the maximum rotation ω_i occurs when P_1 is a critical GLM. In this case, P_1 and Q_1 (the grid point adjacent to P_0 , also rotated by ω_i) both lie on the contour C_1 so that $a^2 y^2 = b^2 (z_1^2 - x^2)$. Hence

$$a^2 r_i^2 \sin^2 \psi_i = b^2 (z_1^2 - r_i^2 \cos^2 \psi_i).$$




Figure 15. Maximum rotation before a GLM becomes critical.

where $r_i^2 = z_0^2 + h^2\|v_i\|^2 - 2h\|v_i\|z_0 \cos \theta_i$ by the cosine rule and $\sin(\psi_i + \omega_i) = h\|v_i\|(\sin \theta_i)/r_i$ by the sine rule.

Since the GLM P_0 becomes critical for rotations ω_i , it follows that equation (8) holds. Furthermore, since the points lie in the right half-plane $z_0 \geq \min\{h\|v_i\| \cos \theta_i\}$. Theorem 4 shows that $z_0 \leq z_{\max}$. \square

Note that although it is the relative rotation of the function and the grid about the minimiser that is important, Figure 15 shows rotation of the grid. This is purely for convenience. Also, only the grid point adjacent to the GLM that becomes critical under the rotation is shown, to avoid clutter.

Figure 16 shows the maximum rotation before a given GLM becomes critical for the example quadratic function and Grid 3 of the sample grids. The data points (marked with asterisks) represent the actual position of the GLMs for this grid. The dashed lines represent $\max\{\omega_i\}$ and the solid lines represents $\omega_{\max} = \min\{\omega_i\}$. Although ω_{\max} is only shown for the example quadratic and Grid 3, the other sample grids all produced strikingly similar graphs.

4. DISCUSSION

Although all grids are equal, as far as the convergence proof for grid-based methods is concerned, the results presented in this paper show that a poor choice of grid can produce GLMs an arbitrary distance from the minimisers of even strictly convex quadratic functions. This has serious implications for practical grid-based algorithms that use only the grid size parameter to determine a suitable stopping criterion. Such algorithms may return arbitrarily bad approximations to stationary points of the function under consideration. Furthermore, an algorithm may rapidly (and prematurely) reduce the grid size if a sequence

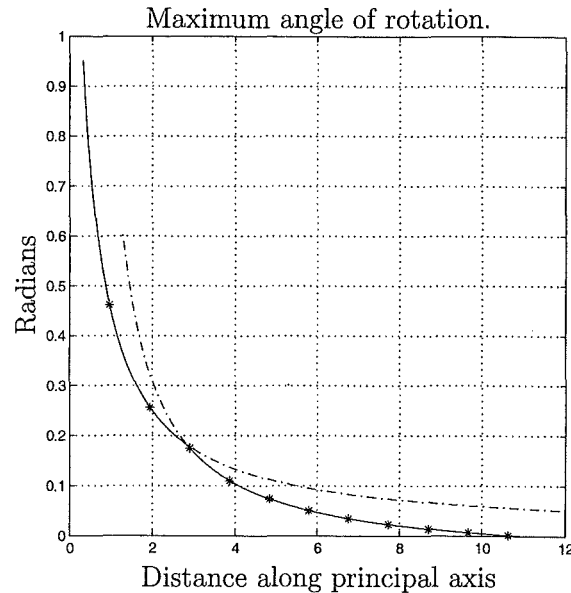


Figure 16. Maximum rotation before a GLM becomes critical for the example quadratic function and Grid 3.

of GLMs is located far from the minimiser. Without an appropriate compensatory strategy, convergence for such algorithms, although guaranteed in theory, may be extremely slow in practice. These problems do not arise with conjugate grids. Conjugate grids guarantee that *any* GLM will be within one cell diameter of the minimiser of a strictly convex quadratic function. Under such circumstances it is appropriate to use the grid size parameter to determine when a practical algorithm could be terminated. Furthermore, subsections 3.1 and 3.2 indicate that conjugate grids are reasonably stable under mild perturbations, so that in practice, grids which are only approximately conjugate are sufficient. The authors of [3] write:

An important aspect of the main convergence result is that successive grids may be arbitrarily translated, rotated, and sheared relative to one another, and each grid axis may be rescaled independently of the others. This flexibility allows second-order information to be incorporated into the shape of successive grids, for example, by aligning grid axes along conjugate directions or the principal axes of an approximating quadratic function.

Hence successively better conjugate approximations could be generated, with the aim of improving an algorithm's practical performance, without affecting its theoretical convergence properties. Similar benefits are expected for non-quadratic functions with grid directions aligned with conjugate directions of an appropriate approximating quadratic function.

Although there is a bound on the smallest angle between pairs of conjugate directions for strictly convex quadratic functions, care must be taken when the Hessian matrix of an approximating quadratic function is singular, or nearly so, at, or near the minimiser. In this situation it is possible for a pair of conjugate directions to become arbitrarily close to being linearly dependent.

Some limited numerical results showing the improved performance of a grid-based method using conjugate grid directions are presented in [1]. Further improvements are currently being investigated.

Finally, we note that the results presented in this paper implicitly make use of a maximal positive basis. These results could be generalised so they apply to grids and GLMs based on any positive basis. However such a generalisation would require an extra layer of notation and theory which, we believe, would obfuscate the ideas presented here.

REFERENCES

- [1] I. D. Coope and C. J. Price. A direct search conjugate directions algorithm for unconstrained minimization. *ANZIAM Journal*, 42(E):C478 – C498, 2000. <http://aniamj.austms.org.au/V42/CTAC99/>.
- [2] I. D. Coope and C. J. Price. Frame based methods for unconstrained optimization. *Journal of Optimization Theory and Applications*, 107(2):261–274, November 2000.
- [3] I. D. Coope and C. J. Price. On the convergence of grid-based methods for unconstrained optimization. *SIAM Journal on Optimization*, 11(4):859–869, 2001.
- [4] I. D. Coope and C. J. Price. Positive bases in numerical optimization. *Computational Optimization and Applications*, 21:169–175, 2002.
- [5] C. Davis. Theory of positive linear dependence. *American Journal of Mathematics*, pages 733–746, 1954.
- [6] E. D. Dolan, R. M. Lewis, and V. Torczon. On the convergence of pattern search, June 2000. Submitted to SIAM Journal on Optimization.
- [7] P. E. Gill, W. Murray, and M. H. Wright. *Practical optimization*. Academic Press, London, 1981.
- [8] R. Hooke and T. A. Jeeves. Direct search solution of numerical and statistical problems. *Journal of the Association for Computing Machinery (ACM)*, 8(2):212–219, April 1961.
- [9] J. Nocedal and S. J. Wright. *Numerical optimization*. Springer, New York, 1999.
- [10] C. J. Price, I. D. Coope, and D. Byatt. A convergent variant of the Nelder-Mead algorithm. *Journal of Optimization Theory and Applications*, 113(1):5–19, April 2002.
- [11] H. H. Rosenbrock. An automatic method for finding the greatest or least value of a function. *The Computer Journal*, 3(3):175–184, October 1960.
- [12] V. Torczon. On the convergence of pattern search algorithms. *SIAM Journal on Optimization*, 7(1):1–25, February 1997.
- [13] M. H. Wright. Direct search methods: once scorned, now respectable. In D. F. Griffiths and G. A. Watson, editors, *Numerical Analysis 1995*, number 344 in Pitman Research Notes in Mathematics Series, pages 191–208, Harlow, UK, 1996. Addison Wesley Longman Ltd.

DEPARTMENT OF MATHEMATICS AND STATISTICS, UNIVERSITY OF CANTERBURY, PRIVATE BAG 4800,
CHRISTCHURCH, NEW ZEALAND

E-mail address: `d.byatt@math.canterbury.ac.nz`

E-mail address: `i.coope@math.canterbury.ac.nz`

E-mail address: `c.price@math.canterbury.ac.nz`