

A projection method for bound-constrained weighted linear least-squares

Working note RAL-NA-2024-1 — Nicholas I. M. Gould

April 10, 2024

1 Introduction

We consider the bound-constrained linear least-squares problem

$$\underset{x \in \mathcal{X}}{\text{minimize}} \quad f(x) := \frac{1}{2} \|Ax - b\|_W^2 \quad \text{where } \mathcal{X} := \{x \in \mathbb{R}^n : x^L \leq x \leq x^U\}. \quad (1.1)$$

Here, we are given a vector of m observations b , a linear model Ax that aims to match b , written in terms of an m by n design matrix A , a positive-definite matrix of weights W , and vectors of lower and upper bounds $x^L, x^U \in \mathbb{R}^n$, some of whose components may be infinite. We are particularly concerned with problems for which n and m are large, and A and W are sparse; in practice W is frequently just a diagonal matrix of strictly positive weights, indeed diagonal weights of one are common.

Throughout we use the Euclidean inner product $\langle \cdot, \cdot \rangle$, the corresponding norm $\|\cdot\| = \|\cdot\|_2$, weighted versions so that $\langle u, v \rangle_W = \langle u, Wv \rangle$ and $\|u\|_W = \sqrt{\langle u, u \rangle_W}$ for any vectors u and v), and the projection operator

$$P_{\mathcal{X}}[x] = \max(x^L, \min(x^U, x)),$$

where \min and \max are applied componentwise. We denote the residual by $r(x) := Ax - b$, and the gradient and (constant) Hessian of f by $g(x) := A^T W r(x)$ and $H := A^T W A$. We partition the indices of $x \in \mathcal{X}$ into bounded and free sets via

$$\mathcal{B}^L(x) := \{j \in \mathbb{N}_n : [x]_j = [x^L]_j\}, \quad \mathcal{B}^U(x) := \{j \in \mathbb{N}_n : [x]_j = [x^U]_j\},$$

$$\mathcal{B}(x) := \mathcal{B}^L(x) \cup \mathcal{B}^U(x) \quad \text{and} \quad \mathcal{F}(x) := \{j \in \mathbb{N}_n : [x^L]_j < [x]_j < [x^U]_j\},$$

where $\mathbb{N}_n := \{1, \dots, n\}$ and $[x]_j$ is the j th component of x . Finally we let $A_{\mathcal{J}}$ be the matrix whose columns are those of A indexed by the set $\mathcal{J} \subseteq \mathbb{N}_n$, I is the n by n identity matrix, $I_{\mathcal{J}}$ be the matrix of columns of I indexed \mathcal{J} , and $e_j \in \mathbb{R}^n$ is the j -th coordinate vector.

2 Method

To solve (1.1), we use an accelerated gradient-projection method. Given a iterate $x_k \in \mathcal{X}$, an improvement x_{k+1} is found as follows:

1. Stop if x_k satisfies suitable termination criteria.
2. Compute

$$x_k^c = P_{\mathcal{X}}[x_k - \alpha_k g(x_k)]$$

for some suitable α_k for which $f(x_k^c) < f(x_k)$.

3. Compute x_k^s as an approximation to

$$\arg \min_{x \in \mathbb{R}^n} f(x) \text{ subject to } x_j = x_l \text{ for } j \in \mathcal{B}^l(x_k^c) \text{ and } x_j = x_u \text{ for } j \in \mathcal{B}^u(x_k^c)$$

4. Compute

$$x_{k+1} = P_{\mathcal{X}}[x_k^c + \beta_k(x_k^s - x_k^c)]$$

for some suitable β_k for which $f(x_{k+1}) \leq f(x_k^c)$.

Steps 3 and 4 may sometimes be omitted if good progress is made in Step 2.

The dominant computations involved are the piecewise linesearches in Steps 2 and 4, and the linear least-squares minimization with fixed variables in Step 3. We consider each in turn.

2.1 Piecewise linesearches

Given a base point x^s and a search direction d , consider the path

$$x(\alpha) := P_{\mathcal{X}}(x^s + \alpha d)$$

for $\alpha \geq 0$. On this path, our aim is to find a point for which $f(x(\alpha))$ is “suitably” smaller than $f(x^s)$. Clearly $x(\alpha)$ is piecewise linear, and changes direction at a finite sequence of “breakpoints” $\alpha_i > \alpha_{i-1}$, for $i = 1, \dots, m$, with $\alpha_0 = 0$. At breakpoint α_i , one or more of the variables $[x]_j$, $j \in \mathcal{B}_i$, encounters a bound, and each is fixed at that value for $\alpha \geq \alpha_i$. In particular,

$$x(\alpha_i + \Delta\alpha) = x_i + \Delta\alpha d_i \tag{2.2}$$

for all $\Delta\alpha \in [0, \Delta\alpha_i]$, where

$$x_i := x(\alpha_i) \equiv P_{\mathcal{X}}(x^s + \alpha_i d),$$

$$[d_i]_j := \begin{cases} 0 & \text{if } j \in \bigcup_{k=0}^i \mathcal{B}_k \text{ or} \\ [d]_j & \text{otherwise,} \end{cases} \tag{2.3}$$

$$\Delta\alpha_i := \alpha_{i+1} - \alpha_i, \tag{2.4}$$

and specifically

$$\mathcal{B}_0 := \{j : [x^S]_j = [x^L]_j \text{ and } [d]_j < 0 \text{ or } [x^S]_j = [x^U]_j \text{ and } [d]_j > 0 \text{ or } [d]_j = 0\}.$$

While it might appear that the breakpoints need be sorted in advance, this is not necessary.

Let

$$\alpha_j^B := \begin{cases} \frac{[x^U]_j - [x^S]_j}{[d]_j} & \text{if } [d]_j > 0, \\ \frac{[x^L]_j - [x^S]_j}{[d]_j} & \text{if } [d]_j < 0 \text{ and} \\ 0 & \text{if } [d]_j = 0 \end{cases} \quad (2.5)$$

for $j \in \mathbb{N}_n$ be the unordered breakpoints. Then the i -th ordered breakpoint may be found efficiently knowing the $i-1$ -st by arranging the $\{\alpha_j^B\}$ in a heap, and using the Heapsort method [10]. In practice, breakpoints that are very close together are considered as a single point.

2.1.1 An exact piecewise minimizer

We now consider $f(x)$ on the segment (2.2). It follows immediately that

$$\begin{aligned} f(x(\alpha_i + \Delta\alpha)) &= f(x_i + \Delta\alpha d_i) = \frac{1}{2} \|A(x_i + \Delta\alpha d_i) - b\|_W^2 \\ &= \frac{1}{2} \|r_i + \Delta\alpha s_i\|_W^2 \\ &= \frac{1}{2} \|r_i\|_W^2 + \Delta\alpha \langle r_i, s_i \rangle_W + \frac{1}{2} \Delta\alpha^2 \|s_i\|_W^2, \end{aligned} \quad (2.6)$$

where

$$r_i := Ax_i - b \quad (2.7)$$

and

$$s_i := Ad_i. \quad (2.8)$$

Notice that (2.6) implies that $f(x(\alpha))$ is convex, and that it achieves its smallest value at

$$\alpha_i^M := \alpha_i - \frac{\langle r_i, s_i \rangle_W}{\|s_i\|_W^2}$$

so long as this lies in the interval $[\alpha_i, \alpha_{i+1}]$. Thus an obvious method to find the global minimum on $x(\alpha)$ is to step through the breakpoints α_i in increasing order until either $\langle r_i, s_i \rangle \geq 0$, in which case the minimizer occurs at $\alpha = \alpha_i$ or $\alpha_i^M \leq \alpha_{i+1}$, in which case a minimizer is at $\alpha = \alpha_i^M$ [2, §3], [3, Alg.17.3.1 with typo corrections]. The process is finite, since there are only at most $n + 1$ breakpoints. As we shall now show, the whole process can be implemented while moving from one segment to the next by evaluating and using the product Av , where v is a vector whose non-zeros occur only in positions corresponding to components of x that reach bounds as the segment ends; typically v is extremely sparse, usually only having a single nonzero.

Superficially, the scheme suggested above requires that we calculate the slope and curvature

$$f'_i := \langle r_i, s_i \rangle_W \text{ and } f''_i := \|s_i\|_W^2,$$

but as we shall now see these may be recurred with modest expense rather than formed afresh. The key is the definition (2.3) of d_i . In particular

$$d_{i+1} = d_i - v_{i+1}, \quad (2.9)$$

where

$$v_{i+1} := \sum_{j \in \mathcal{B}_{i+1}} [d]_j e_j,$$

and v_{i+1} is almost certainly a very sparse vector. Thus since A is sparse, it is also highly likely that so are

$$p_{i+1} := Av_{i+1} \quad \text{and} \quad y_{i+1} := Wp_{i+1}; \quad (2.10)$$

each column of A can only be accessed a single time during the entire iteration. Hence, using (2.8), (2.9) and (2.10),

$$s_{i+1} = s_i - p_{i+1} \quad (2.11)$$

differs from s_i in only a few components and therefore may be updated efficiently rather than recomputed. Therefore

$$\begin{aligned} f''_{i+1} &= \|s_{i+1}\|_W^2 = \|s_i - p_{i+1}\|_W^2 = \|s_i\|_W^2 + \langle p_{i+1} - 2s_i, p_{i+1} \rangle_W \\ &= f''_i + \langle p_{i+1} - 2s_i, p_{i+1} \rangle_W \\ &= f''_i + \langle p_{i+1} - 2s_i, y_{i+1} \rangle \end{aligned}$$

which may be updated using a sparse inner product involving y_{i+1} .

Recurring the slope $f'_i = \langle r_i, s_i \rangle_W$ is only slightly more awkward. To proceed, let

$$r^S := Ax^S - b, \quad (2.12)$$

$$g^S := A^T W r^S, \quad (2.13)$$

$$\Delta x_i := x_i - x^S \quad \text{and} \quad (2.14)$$

$$q_i := A \Delta x_i. \quad (2.15)$$

It then follows from (2.7), (2.12), (2.14) and (2.15) that

$$r_i = A(x^S + \Delta x_i) - b = Ax^S - b + A \Delta x_i = Ax^S - b + q_i = r^S + q_i,$$

and hence from (2.8) and (2.13) that

$$\langle r_i, s_i \rangle_W = \langle r^S + q_i, s_i \rangle_W = \langle r^S, s_i \rangle_W + \langle q_i, s_i \rangle_W = l_i + \gamma_i,$$

where

$$l_i := \langle g^S, d_i \rangle \quad \text{and} \quad \gamma_i := \langle q_i, s_i \rangle_W. \quad (2.16)$$

Thus (2.9), (2.10) and (2.16) give

$$\begin{aligned} l_{i+1} &= \langle g^S, d_{i+1} \rangle = \langle g^S, d_i - v_{i+1} \rangle = \langle g^S, d_i \rangle - \langle g^S, v_{i+1} \rangle \\ &= l_i - \langle g^S, v_{i+1} \rangle = l_i - \langle r^S, WAv_{i+1} \rangle = l_i - \langle r^S, y_{i+1} \rangle, \end{aligned}$$

and so l_{i+1} may be updated from l_i using a sparse inner product involving either v_{i+1} or y_{i+1} . Since (2.2) gives

$$x_{i+1} = x_i + \Delta\alpha_i d_i,$$

it follows from (2.8), (2.14) and (2.15) that

$$\begin{aligned} q_{i+1} &= A\Delta x_{i+1} = A(x_{i+1} - x^s) = A(x_i + \Delta\alpha_i d_i - x^s) \\ &= A(x_i - x^s) + \Delta\alpha_i A d_i = A\Delta x_i + \Delta\alpha_i A d_i \\ &= q_i + \Delta\alpha_i s_i. \end{aligned} \tag{2.17}$$

Combining this with (2.11), we find

$$\begin{aligned} \gamma_{i+1} &= \langle q_{i+1}, s_{i+1} \rangle_W = \langle q_i + \Delta\alpha_i s_i, s_i - p_{i+1} \rangle_W \\ &= \langle q_i, s_i \rangle_W + \Delta\alpha_i \langle s_i, s_i \rangle_W - \langle q_{i+1}, p_{i+1} \rangle_W \\ &= \gamma_i + \Delta\alpha_i f_i'' - \langle q_{i+1}, p_{i+1} \rangle_W, \\ &= \gamma_i + \Delta\alpha_i f_i'' - \langle q_{i+1}, y_{i+1} \rangle, \end{aligned} \tag{2.18}$$

where we only need to take the inner product over components j for which $[y_{i+1}]_j \neq 0$. Unfortunately, the update (2.17) for q_{i+1} does not involve a sparse vector, and we need to dig a little deeper. The secret is to define

$$u_{i+1} := q_{i+1} - \alpha_{i+1} s_i, \quad \text{with } u_0 = 0. \tag{2.19}$$

Then it follows from (2.17), $\Delta\alpha_i - \alpha_{i+1} = -\alpha_i$ from (2.4), and (2.11) that

$$\begin{aligned} u_{i+1} &= q_i + \Delta\alpha_i s_i - \alpha_{i+1} s_i = q_i + (\Delta\alpha_i - \alpha_{i+1}) s_i = q_i - \alpha_i s_i = q_i - \alpha_i (s_{i-1} - p_i) \\ &= u_i + \alpha_i p_i, \end{aligned}$$

and this is a sparse update. Thus rather than recurring q_i , we may instead recur u_i and obtain $q_{i+1} = u_{i+1} + \alpha_{i+1} s_i$ from (2.19) as required. The important difference is that the recursions for u_{i+1} and s_i only involve the likely-sparse p_i . In particular, the recurrence for γ_{i+1} in (2.18) becomes

$$\begin{aligned} \gamma_{i+1} &= \gamma_i + \Delta\alpha_i f_i'' - \langle u_{i+1}, p_{i+1} \rangle_W - \alpha_{i+1} \langle s_i, p_{i+1} \rangle_W \\ &= \gamma_i + \Delta\alpha_i f_i'' - \langle u_{i+1}, y_{i+1} \rangle - \alpha_{i+1} \langle s_i, y_{i+1} \rangle. \end{aligned}$$

We summarize our findings as Algorithm 2.1; this is essentially a specific case of a more general framework [5, Alg.4]. Notice that we also record the value $f_i = f(x_i)$ as a bi-product, and that this may be updated using (2.6) as we proceed from breakpoint to breakpoint.

Algorithm 2.1: Finding the piecewise arc minimizer $x^c = P_{\mathcal{X}}(x^s + \alpha^c d)$ of f

0. Initialization: The initial point $x^s \in \mathcal{X}$ and search direction d are given. Compute the residual $r^s = Ax^s - b$, and the breakpoints α_j^B from (2.5) for all $j \in \mathbb{N}_n$.

Let $\alpha_0 = 0$, $u_0 = 0$,

$$\mathcal{B}_0 = \{j \in \mathbb{N}_n : \alpha_j^{\text{B}} = 0\} \text{ and } d_0 = d - \sum_{j \in \mathcal{B}_0} [d]_j e_j,$$

compute

$$\begin{aligned} p_0 &= Ad_0 \text{ and } y_0 = Wp_0, \\ f_0 &= \frac{1}{2} \langle r^{\text{S}}, Wr^{\text{S}} \rangle, \quad f'_0 = \langle r^{\text{S}}, y_0 \rangle \text{ and } f''_0 = \langle p_0, y_0 \rangle \end{aligned}$$

and set $s_0 = p_0$ and $i = 0$.

1. Find the next breakpoint: Determine α_{i+1} , the first breakpoint beyond α_i .

2. Check the current interval for arc minimizer:

If $f'_i \geq 0$, set $\alpha^{\text{C}} = \alpha_i$, $x^{\text{C}} = P_{\mathcal{X}}[x^{\text{S}} + \alpha^{\text{C}}d]$, $f(x^{\text{C}}) = f_i$, and stop.

If $f''_i > 0$ and $\alpha_i - f'_i/f''_i \leq \alpha_{i+1}$, set $\alpha^{\text{C}} = \alpha_i - f'_i/f''_i$, $x^{\text{C}} = P_{\mathcal{X}}[x^{\text{S}} + \alpha^{\text{C}}d]$
 $f(x^{\text{C}}) = f_i + (\alpha^{\text{C}} - \alpha_i)f'_i + \frac{1}{2}(\alpha^{\text{C}} - \alpha_i)^2 f''_i$, and stop.

3. Prepare for the next interval: Set $\Delta\alpha_i = \alpha_{i+1} - \alpha_i$, recur

$$u_{i+1} = u_i + \alpha_i p_i,$$

let

$$\mathcal{B}_{i+1} = \{j \in \mathbb{N}_n : \alpha_j^{\text{B}} = \alpha_{i+1}\} \text{ and } v_{i+1} = \sum_{j \in \mathcal{B}_{i+1}} [d]_j e_j,$$

and compute

$$p_{i+1} = Av_{i+1} \text{ and } y_{i+1} = Wp_{i+1}.$$

4. Compute the value, slope and curvature: Compute

$$\begin{aligned} f_{i+1} &= f_i + \Delta\alpha_i f'_i + \frac{1}{2}(\Delta\alpha_i)^2 f''_i, \\ f'_{i+1} &= f'_i + \Delta\alpha_i f''_i - \langle r^{\text{S}} + u_{i+1} + \alpha_{i+1} s_i, y_{i+1} \rangle \text{ and} \\ f''_{i+1} &= f''_i + \langle p_{i+1} - 2s_i, y_{i+1} \rangle, \end{aligned}$$

update

$$s_{i+1} = s_i - p_{i+1},$$

increment i by 1 and return to Step 1.

In practice, we compute f'_{i+1} afresh when $|f'_{i+1}/f'_i|$ becomes small to guard against possible accumulated rounding errors in the recurrences. An earlier variant [1] based on algorithms for general quadratic objectives [2, §3], but specialised for the least-squares case, required products with both A and A^T at each breakpoint.

2.1.2 An approximate piecewise minimizer

In some cases, it may instead be advantageous to approximate the piecewise minimizer using a safeguarded, backtracking, piecewise linesearch [8]. The idea is simply to pick an initial stepsize α_0 , backtracking reduction factor $\beta \in (0, 1)$ and decrease tolerance $\eta \in (0, \frac{1}{2})$, and to choose a sequence of decreasing stepsizes $\alpha_i := \alpha_0 \beta^i$ for increasing $i \geq 0$ until

$$f(x_i) \leq f(x^s) + \eta \langle d_i, g^s \rangle, \quad (2.20)$$

involving the trial point $x_i = P_{\mathcal{X}}[x^s + \alpha_i d]$, the trial step $d_i = x_i - x^s$, and the gradient at the base point $g^s = A^T W(Ax^s - b)$.

To do so, we take advantage of the structure of the trial step d_i and basic properties of the backtracking projected line search. In particular, we know that once a component, say the j th, satisfies $[x_i^L]_j < [x_i]_j < [x_i^U]_j$, then $[x_\ell^L]_j < [x_\ell]_j < [x_\ell^U]_j$ will continue to hold for all $\ell \geq i$. Thus, by contrast to the search for the exact minimizer in Section 2.1.1 that moves forward along the piecewise projected gradient path fixing variables, the search here frees variables from their bounds as it proceeds backwards towards x^s .

With this in mind, we compute the index set

$$\mathcal{A}^s = \{j : [d]_j = 0 \text{ or } [x^s]_j = [x^L]_j \text{ and } [d]_j < 0 \text{ or } [x^s]_j = [x^U]_j \text{ and } [d]_j > 0\} \quad (2.21)$$

of components that are fixed at the base point on the piecewise search arc, record the vector x^B for which

$$[x^B]_j := \begin{cases} [x^s]_j & \text{if } j \in \mathcal{A}^s, \\ [x^L]_j & \text{if } j \notin \mathcal{A}^s \text{ and } [d]_j < 0, \text{ and} \\ [x^U]_j & \text{if } j \notin \mathcal{A}^s \text{ and } [d]_j > 0, \end{cases} \quad (2.22)$$

and maintain the sets of active (i.e., fixed) and free components

$$\mathcal{A}_i := \{j \notin \mathcal{A}^s : [x_i]_j = [x^B]_j\} \text{ and} \quad (2.23)$$

$$\mathcal{F}_i := \{j \notin \mathcal{A}^s : [x_i]_j \neq [x^B]_j\}, \quad (2.24)$$

for $i \geq 0$, as well as the intermediate components, those that change from active to free,

$$\mathcal{I}_i := \mathcal{F}_i \cap \mathcal{A}_{i-1}, \quad (2.25)$$

for $i \geq 1$, at x_i . As we have already mentioned,

$$\mathcal{F}_i \subseteq \mathcal{F}_{i+1} \text{ and } \mathcal{A}_{i+1} \subseteq \mathcal{A}_i \text{ for all } i \geq 0$$

as a consequence of the approximate piecewise search.

It is immediate that, permuting the variables in the obvious way,

$$d_i = \begin{pmatrix} s_{\mathcal{A}_i} \\ \alpha_i d_{\mathcal{F}_i} \end{pmatrix}, \quad (2.26)$$

where

$$s := x^B - x^S.$$

Hence

$$r_i := Ax_i - b = A(x^S + d_i) - b = r^S + Ad_i = r^S + A_{\mathcal{A}_i}s_{\mathcal{A}_i} + \alpha_i A_{\mathcal{F}_i}d_{\mathcal{F}_i} = r_i^A + \alpha_i r_i^F,$$

where

$$r_i^A := r^S + A_{\mathcal{A}_i}s_{\mathcal{A}_i} \quad \text{and} \quad r_i^F := A_{\mathcal{F}_i}d_{\mathcal{F}_i}, \quad (2.27)$$

and therefore

$$f(x_i) = \frac{1}{2} \|r_i^A\|_W^2 + \alpha_i \langle r_i^A, r_i^F \rangle_W + \frac{1}{2} \alpha_i^2 \|r_i^F\|_W^2.$$

Thus $f(x_i)$ may be found trivially from the three quantities

$$f_i^C := \frac{1}{2} \|r_i^A\|_W^2, \quad f_i^L := \langle r_i^A, r_i^F \rangle_W \quad \text{and} \quad f_i^Q := \frac{1}{2} \|r_i^F\|_W^2. \quad (2.28)$$

Although it is possible to compute these quantities afresh, it is usually more efficient to recur them as the piecewise linesearch proceeds instead. To do so, note that

$$\mathcal{A}_i = \mathcal{A}_{i-1} \setminus \mathcal{I}_i \quad \text{and} \quad \mathcal{F}_i = \mathcal{F}_{i-1} \cup \mathcal{I}_i, \quad (2.29)$$

and thus that

$$r_i^A = r_{i-1}^A - p_i \quad \text{and} \quad r_i^F = r_{i-1}^F + q_i, \quad (2.30)$$

where

$$p_i := A_{\mathcal{I}_i}s_{\mathcal{I}_i} \quad \text{and} \quad q_i := A_{\mathcal{I}_i}d_{\mathcal{I}_i}. \quad (2.31)$$

Therefore, when A is sparse and if a modest number of variables move off bounds at x_i , the vectors p_i and q_i will most likely be sparse, indeed, aside from exact cancellations, their nonzeros will occur in the same positions. It then follows from (2.28) and (2.31) that

$$\begin{aligned} f_i^C &= f_{i-1}^C - \langle p_i, r_{i-1}^A \rangle_W + \frac{1}{2} \|p_i\|_W^2 \\ &= f_{i-1}^C - \langle y_i, r_{i-1}^A - \frac{1}{2}p_i \rangle, \\ f_i^L &= f_{i-1}^L + \langle q_i, r_{i-1}^A \rangle_W - \langle p_i, r_{i-1}^F \rangle_W - \langle q_i, p_i \rangle_W \\ &= f_{i-1}^L + \langle z_i, r_{i-1}^A \rangle - \langle y_i, r_{i-1}^F \rangle - \frac{1}{2} \langle z_i, p_i \rangle - \frac{1}{2} \langle y_i, q_i \rangle \\ &= f_{i-1}^L + \langle z_i, r_{i-1}^A - \frac{1}{2}p_i \rangle - \langle y_i, r_{i-1}^F + \frac{1}{2}q_i \rangle \quad \text{and} \\ f_i^Q &= f_{i-1}^Q + \langle q_i, r_{i-1}^F \rangle_W + \frac{1}{2} \|q_i\|_W^2 \\ &= f_{i-1}^Q + \langle z_i, r_{i-1}^F + \frac{1}{2}q_i \rangle, \end{aligned} \quad (2.32)$$

involving inner products with the likely-sparse vectors

$$y_i := Wp_i \quad \text{and} \quad z_i := Wq_i. \quad (2.33)$$

In order to check (2.20), we also need to compute

$$\langle d_i, g^S \rangle = \langle A^T d_i, r^S \rangle_W = \langle A_{\mathcal{A}_i}s_{\mathcal{A}_i}, r^S \rangle_W + \alpha_i \langle A_{\mathcal{F}_i}d_{\mathcal{F}_i}, r^S \rangle_W = \gamma_i^A + \alpha_i \gamma_i^F, \quad (2.34)$$

where

$$\gamma_i^A := \langle A_{\mathcal{A}_i} s_{\mathcal{A}_i}, r^S \rangle_W \quad \text{and} \quad \gamma_i^F := \langle A_{\mathcal{F}_i} d_{\mathcal{F}_i}, r^S \rangle_W,$$

using (2.26). It then follows from (2.29) and (2.31) that

$$\begin{aligned} \gamma_i^A &= \langle A_{\mathcal{A}_{i-1}} s_{\mathcal{A}_{i-1}}, r^S \rangle_W - \langle A_{\mathcal{I}_i} s_{\mathcal{I}_i}, r^S \rangle_W = \gamma_{i-1}^A - \langle p_i, r^S \rangle_W \\ &= \gamma_{i-1}^A - \langle y_i, r^S \rangle \quad \text{and} \\ \gamma_i^F &= \langle A_{\mathcal{F}_{i-1}} d_{\mathcal{F}_{i-1}}, r^S \rangle_W + \langle A_{\mathcal{I}_i} d_{\mathcal{I}_i}, r^S \rangle_W = \gamma_{i-1}^F + \langle q_i, r^S \rangle_W \\ &= \gamma_{i-1}^F + \langle z_i, r^S \rangle. \end{aligned} \tag{2.35}$$

Thus we may obtain (2.34) by recurring γ_i^A and γ_i^F using (2.35), and the latter simply requires a further pair of sparse inner products.

It is important to notice that the crucial likely-sparse vectors p_i and q_i in (2.31) needed by the recurrences (2.32) and (2.35) only depend on the set of indices \mathcal{I}_i that change status from fixed to free during the i -th backtrack. Although formally we define this using x_i , in practice we do not need to form x_i , indeed to do so would require a projection $P_{\mathcal{X}}[x^S + \alpha_i d]$ for each attempted step α_i . Instead, just as in Section 2.1.1, we find the unordered breakpoints (2.5), and then arrange them in a “backward” heap starting at the first for which $\alpha_j^B > \alpha_0$. We then adjust the heap to find precisely the indices I_{i+1} of those between α_i and α_{i+1} using the Heapsort algorithm as required.

We summarize our discussion as Algorithm 2.2.

Algorithm 2.2: Find an approximate backtracking arc minimizer x^C of f

0. Initialization: The initial point $x^S \in \mathcal{X}$, search direction d , initial stepsize $\alpha_0 > 0$, reduction factor $\beta \in (0, 1)$ and decrease tolerance $\eta \in (0, \frac{1}{2})$ are given. Compute the residual $r^S = Ax^S - b$, the initial objective value $f(x^S) = \frac{1}{2} \|r^S\|_W^2$, the base fixed set \mathcal{A}^S from (2.21), the breakpoints α_j^B from (2.5) for all $j \in \mathbb{N}_n$, the end of the arc x^B from (2.22) and its direction $s = x^B - x^S$, the initial search point $x_0 = P_{\mathcal{X}}[x^S + \alpha_0 d]$, the active and free components at x_0 ,

$$\mathcal{A}_0 = \{j : \alpha_0 > \alpha_j^B\} \quad \text{and} \quad \mathcal{F}_0 = \{j \notin \mathcal{A}^S : \alpha_0 \leq \alpha_j^B\},$$

and the corresponding residuals

$$r_0^A = r^S + p_0 \quad \text{and} \quad r_0^F = q_0$$

using the matrix-vector products

$$p_0 = A_{\mathcal{A}_0} s_{\mathcal{A}_0} \quad \text{and} \quad q_0 = A_{\mathcal{F}_0} d_{\mathcal{F}_0}.$$

Initialize

$$f_0^C = \frac{1}{2} \|r_0^A\|_W^2, \quad f_0^L = \langle r_0^A, r_0^F \rangle_W \quad \text{and} \quad f_0^Q = \frac{1}{2} \|r_0^F\|_W^2,$$

as well as

$$\gamma_0^A = \langle p_0, r^S \rangle_W \quad \text{and} \quad \gamma_0^F = \langle q_0, r^S \rangle_W.$$

Set $i = 0$.

1. Check for an approximate arc minimizer: Compute

$$f_i = f_i^C + \alpha_i f_i^L + \alpha_i^2 f_i^Q \quad \text{and} \quad \gamma_i = \gamma_i^A + \alpha_i \gamma_i^F.$$

If $f_i \leq f(x^S) + \eta \gamma_i$, set

$$\alpha^C = \alpha_i, \quad x^C = P_{\mathcal{X}}[x^S + \alpha^C d] \quad \text{and} \quad f(x^C) = f_i,$$

and stop.

2. Find the next set of indices that change status: Let $\alpha_{i+1} = \beta \alpha_i$, and compute

$$\mathcal{I}_{i+1} = \{j : \alpha_{i+1} < \alpha_j^B \leq \alpha_i\}$$

using the Heapsort algorithm.

3. Update the components of the objective and its slope: Compute

$$\begin{aligned} p_{i+1} &= A_{\mathcal{I}_{i+1}} s_{\mathcal{I}_{i+1}}, \quad y_{i+1} = W p_{i+1}, \\ q_{i+1} &= A_{\mathcal{I}_{i+1}} d_{\mathcal{I}_{i+1}} \quad \text{and} \quad z_{i+1} = W q_{i+1}, \end{aligned}$$

update

$$\begin{aligned} f_{i+1}^C &= f_i^C - \langle y_{i+1}, r_i^A - \tfrac{1}{2} p_{i+1} \rangle, \\ f_{i+1}^L &= f_i^L + \langle z_{i+1}, r_i^A - \tfrac{1}{2} p_{i+1} \rangle - \langle y_{i+1}, r_i^F + \tfrac{1}{2} q_{i+1} \rangle, \\ f_{i+1}^Q &= f_i^Q + \langle z_{i+1}, r_i^F + \tfrac{1}{2} q_{i+1} \rangle, \\ \gamma_{i+1}^A &= \gamma_i^A - \langle y_{i+1}, r^S \rangle, \\ \gamma_{i+1}^F &= \gamma_i^F + \langle z_{i+1}, r^S \rangle, \\ r_{i+1}^A &= r_i^A - p_{i+1} \quad \text{and} \\ r_{i+1}^F &= r_i^F + q_{i+1}, \end{aligned}$$

increment i by 1 and return to Step 1.

It is also possible to contemplate a variant in which the iterates advances further along the piecewise arc if the initial point x_0 is acceptable. To be specific, if

$$f(x_0) \leq f(x^S) + \eta \langle x_0 - x^S, g^S \rangle, \quad (2.36)$$

we terminate at the arc point x_i with the smallest $i \geq 0$ for which

$$f(x_{i+1}) > f(x^S) + \eta \langle x_{i+1} - x^S, g^S \rangle, \quad (2.37)$$

where, as before, $x_i = P_{\mathcal{X}}[x^S + \alpha_i d]$, but now $\alpha_i = \alpha_0 \beta^{-i} \geq \alpha_0$. The only essential difference is that in this case

$$\mathcal{F}_{i+1} \subseteq \mathcal{F}_i \quad \text{and} \quad \mathcal{A}_i \subseteq \mathcal{A}_{i+1} \quad \text{for all } i \geq 0,$$

and components in $\mathcal{J}_{i+1} := \mathcal{A}_{i+1} \cap \mathcal{F}_i$ change from free to active, i.e.,

$$\mathcal{A}_{i+1} = \mathcal{A}_i \cup \mathcal{J}_{i+1} \quad \text{and} \quad \mathcal{F}_{i+1} = \mathcal{F}_i \setminus \mathcal{J}_{i+1}. \quad (2.38)$$

This leads to

$$r_{i+1}^A = r_i^A + p_{i+1} \quad \text{and} \quad r_{i+1}^F = r_i^F - q_{i+1},$$

where p_i and q_i are given by (2.31), and the obvious changes to (2.32) and (2.35). We summarize the necessary enhancements in Algorithm 2.3.

Algorithm 2.3: Find an approximate piecewise arc minimizer x^C of f

The same as Algorithm 2.2 on page 9 except that Step 1 becomes

1. Check for an approximate backtracking arc minimizer: Compute

$$f_i = f_i^C + \alpha_i f_i^L + \alpha_i^2 f_i^Q \quad \text{and} \quad \gamma_i = \gamma_i^A + \alpha_i \gamma_i^F.$$

If $f_i \leq f(x^S) + \eta \gamma_i$, go to Step 4 if $i = 0$ but otherwise, i.e., if $i > 0$, set

$$\alpha^C = \alpha_i, \quad x^C = P_{\mathcal{X}}[x^S + \alpha^C d] \quad \text{and} \quad f(x^C) = f_i,$$

and stop.

and additionally

4. Find the next set of indices that change status: Let $\alpha_{i+1} = \beta^{-1} \alpha_i$, and compute

$$\mathcal{J}_{i+1} = \{j : \alpha_i < \alpha_j^B \leq \alpha_{i+1}\}$$

using the Heapsort algorithm.

5. Update the components of the objective and its slope: Compute

$$p_{i+1} = A_{\mathcal{J}_{i+1}} s_{\mathcal{J}_{i+1}}, \quad y_{i+1} = W p_{i+1},$$

$$q_{i+1} = A_{\mathcal{J}_{i+1}} d_{\mathcal{J}_{i+1}} \quad \text{and} \quad z_{i+1} = W q_{i+1},$$

update

$$f_{i+1}^C = f_i^C + \langle y_{i+1}, r_i^A + \frac{1}{2} p_{i+1} \rangle,$$

$$f_{i+1}^L = f_i^L - \langle z_{i+1}, r_i^A + \frac{1}{2} p_{i+1} \rangle + \langle y_{i+1}, r_i^F - \frac{1}{2} q_{i+1} \rangle,$$

$$f_{i+1}^Q = f_i^Q - \langle z_{i+1}, r_i^F - \frac{1}{2} q_{i+1} \rangle,$$

$$\gamma_{i+1}^A = \gamma_i^A + \langle y_{i+1}, r^S \rangle,$$

$$\gamma_{i+1}^F = \gamma_i^F - \langle z_{i+1}, r^S \rangle,$$

$$r_{i+1}^A = r_i^A + p_{i+1} \quad \text{and}$$

$$r_{i+1}^F = r_i^F - q_{i+1}.$$

6. Check for an approximate extended arc minimizer: Compute

$$f_{i+1} = f_{i+1}^C + \alpha_{i+1} f_{i+1}^L + \alpha_{i+1}^2 f_{i+1}^Q \quad \text{and} \quad \gamma_{i+1} = \gamma_{i+1}^A + \alpha_{i+1} \gamma_{i+1}^F.$$

If $f_{i+1} > f(x^S) + \eta \gamma_{i+1}$ or $\alpha_{i+1} \geq \max_j \alpha_j^B$, set

$$\alpha^C = \alpha_i, \quad x^C = P_{\mathcal{X}}[x^S + \alpha^C d] \quad \text{and} \quad f(x^C) = f_i,$$

and stop. Otherwise, increment i by 1 and return to Step 4.

Notice the extra stopping check in Step 6: this is simply to prevent the search extending beyond the end of the piecewise arc.

2.2 Linear least-squares minimization with fixed variables

Let $Z[v]$ be the operator that sets the components of a given vector v to zero, i.e.,

$$[Z[v]]_i = \begin{cases} 0 & i \in \mathcal{Z} \\ [v]_i & \text{otherwise,} \end{cases}$$

for a specified subset $\mathcal{Z} \subseteq \mathbb{N}_n$ of indices of x that are to be fixed; in the context of Step 3 in the generic framework described at the start of Section 2, \mathcal{Z} at iteration k will be $\mathcal{B}^L(x_k^c) \cup \mathcal{B}^U(x_k^c)$. To minimize $f(x)$ over the set of variables that are in $\mathcal{F} := \mathbb{N}_n \setminus \mathcal{Z}$, while fixing the remaining components at the values that they have at x_0 , we may apply the following well-known preconditioned conjugate-gradient iterative scheme [7, 9]—here the preconditioner P may be any symmetric, positive-definite matrix, for which the cost of solving $Pv = g$ is modest.

Algorithm 2.4: The preconditioned conjugate-gradient least-squares method

Given x_0 , set $r_0 = Ax_0 - b$ and $g_0 = Z[A^T W r_0]$, and let $v_0 = Z[P^{-1}g_0]$ and $p_0 = -v_0$. For $k = 0, 1, \dots$ until convergence, perform the iteration

$$\begin{aligned} q_k &= Ap_k, \\ \alpha_k &= \langle g_k, v_k \rangle / \langle q_k, W q_k \rangle, \\ x_{k+1} &= x_k + \alpha_k p_k, \\ r_{k+1} &= r_k + \alpha_k q_k, \\ g_{k+1} &= Z[A^T W r_{k+1}], \\ v_{k+1} &= Z[P^{-1}g_{k+1}], \\ \beta_k &= \langle g_{k+1}, v_{k+1} \rangle / \langle g_k, v_k \rangle \text{ and} \\ p_{k+1} &= -v_{k+1} + \beta_k p_k. \end{aligned}$$

Notice that the product $q_k = Ap_k$ only requires access to the columns of A with indices that lie in \mathcal{F} . Likewise, only the components of $A^T W r_{k+1}$ with indices that lie in \mathcal{F} are needed. A good “preconditioner” P will be such that the eigenvalues of $P^{-1}A_{\mathcal{F}}^T W A_{\mathcal{F}}$ are clustered around one, although it is not necessarily easy to achieve this [6]; at the very least, picking $P = \text{diag}(A^T W A)$ is often helpful, and in particular the required diagonal entries are then simply the squares of the W -norms of the columns of A .

3 Regularization

In practice, it is common to consider the regularized bound-constrained linear least-squares problem

$$\underset{x \in \mathcal{X}}{\text{minimize}} \quad \phi(x, \sigma) := \frac{1}{2} \|Ax - b\|_W^2 + \frac{1}{2} \sigma \|x\|_R^2, \quad (3.39)$$

in which we allow an extra regularization term in the positive-definite R -norm, with regularization weight $\sigma \geq 0$; while it is not necessary, in what follows we shall assume that R is sparse so that products between R and sparse vectors remain sparse. Since $\phi(x, \sigma)$ is a linear sum of $f(x)$ and the regularization term, we may use linearity to compute derivatives involving those that we have already seen plus σ times those of $\rho(x) := \frac{1}{2} \|x\|_R^2$. Trivially, the gradient and Hessian of ρ are Rx and R respectively. Equally $\rho(x)$ may be interpreted as a weighted sum-of-squares function $\frac{1}{2} \|Ax - b\|_W^2$ in the special case for which $A = I$, $b = 0$ and $W = R$. Such a perspective then simply allows us to derive the necessary changes in $\rho(x)$ as we investigate the piecewise-linear path $x(\alpha)$.

To generalise the method for finding the exact piecewise minimizer described in Section 2.1.1, we must consider $\rho(x)$ on the segment (2.2). Plainly we have that

$$\rho(x(\alpha_i + \Delta\alpha)) = \rho(x_i + \Delta\alpha d_i) = \frac{1}{2} \|x_i + \Delta\alpha d_i\|_R^2 = \rho_i + \Delta\alpha \rho'_i + \frac{1}{2} \Delta\alpha^2 \rho''_i \quad (3.40)$$

involving the value, slope and curvature

$$\rho_i := \frac{1}{2} \|x_i\|_R^2, \quad \rho'_i := \langle x_i, d_i \rangle_R \quad \text{and} \quad \rho''_i := \|d_i\|_R^2$$

at breakpoint i . Mechanically repeating the arguments in Section 2.1.1, leads to the following generalisation of Algorithm 2.1 for the regularization case.

Algorithm 3.1: Finding the piecewise arc minimizer $x^c = P_{\mathcal{X}}(x^s + \alpha^c d)$ of ϕ

0. Initialization: The initial point $x^s \in \mathcal{X}$ and search direction d are given. Compute the residual $r^s = Ax^s - b$, and the breakpoints α_j^B from (2.5) for all $j \in \mathbb{N}_n$. Let $\alpha_0 = 0$, $u_0 = 0$, $w_0 = 0$,

$$\mathcal{B}_0 = \{j \in \mathbb{N}_n : \alpha_j^B = 0\} \quad \text{and} \quad d_0 = d - \sum_{j \in \mathcal{B}_0} [d]_j e_j,$$

compute

$$\begin{aligned} p_0 &= Ad_0, \quad y_0 = Wp_0 \quad \text{and} \quad z_0 = Rd_0, \\ f_0 &= \frac{1}{2} \langle r^s, W r^s \rangle, \quad f'_0 = \langle r^s, y_0 \rangle \quad \text{and} \quad f''_0 = \langle p_0, y_0 \rangle, \\ \rho_0 &= \frac{1}{2} \langle x^s, R x^s \rangle, \quad \rho'_0 = \langle x^s, z_0 \rangle \quad \text{and} \quad \rho''_0 = \langle d_0, z_0 \rangle, \end{aligned}$$

and set $v_0 = d_0$, $s_0 = p_0$ and $i = 0$.

- 1. Find the next breakpoint:** Determine α_{i+1} , the first breakpoint beyond α_i .
- 2. Check the current interval for arc minimizer:**

Compute

$$\phi_i = f_i + \sigma \rho_i, \quad \phi'_i = f'_i + \sigma \rho'_i \quad \text{and} \quad \phi''_i = f''_i + \sigma \rho''_i.$$

If $\phi'_i \geq 0$, set $\alpha^c = \alpha_i$, $x^c = P_{\mathcal{X}}[x^s + \alpha^c d]$, $\phi(x^c, \sigma) = \phi_i$, and stop.

If $\phi''_i > 0$ and $\alpha_i - \phi'_i/\phi''_i \leq \alpha_{i+1}$, set $\alpha^c = \alpha_i - \phi'_i/\phi''_i$, $x^c = P_{\mathcal{X}}[x^s + \alpha^c d]$
 $\phi(x^c, \sigma) = \phi_i + (\alpha^c - \alpha_i)\phi'_i + \frac{1}{2}(\alpha^c - \alpha_i)^2\phi''_i$, and stop.

3. Prepare for the next interval: Set $\Delta\alpha_i = \alpha_{i+1} - \alpha_i$, recur

$$w_{i+1} = w_i + \alpha_i v_i,$$

$$u_{i+1} = u_i + \alpha_i p_i,$$

let

$$\mathcal{B}_{i+1} = \{j \in \mathbb{N}_n : \alpha_j^B = \alpha_{i+1}\} \quad \text{and} \quad v_{i+1} = \sum_{j \in \mathcal{B}_{i+1}} [d]_j e_j,$$

and compute

$$p_{i+1} = A v_{i+1}, \quad y_{i+1} = W p_{i+1} \quad \text{and} \quad z_{i+1} = R v_{i+1}.$$

4. Compute the value, slope and curvature: Compute

$$f_{i+1} = f_i + \Delta\alpha_i f'_i + \frac{1}{2}(\Delta\alpha_i)^2 f''_i,$$

$$f'_{i+1} = f'_i + \Delta\alpha_i f''_i - \langle r^s + u_{i+1} + \alpha_{i+1} s_i, y_{i+1} \rangle,$$

$$f''_{i+1} = f''_i + \langle p_{i+1} - 2s_i, y_{i+1} \rangle,$$

$$\rho_{i+1} = \rho_i + \Delta\alpha_i \rho'_i + \frac{1}{2}(\Delta\alpha_i)^2 \rho''_i,$$

$$\rho'_{i+1} = \rho'_i + \Delta\alpha_i \rho''_i - \langle x^s + w_{i+1} + \alpha_{i+1} d_i, z_{i+1} \rangle \quad \text{and} \quad (3.41)$$

$$\rho''_{i+1} = \rho''_i + \langle v_{i+1} - 2d_i, z_{i+1} \rangle, \quad (3.42)$$

update

$$d_{i+1} = d_i - v_{i+1},$$

$$s_{i+1} = s_i - p_{i+1},$$

increment i by 1 and return to Step 1.

Slightly less obviously, it is straightforward to show that

$$\langle d_i, z_{i+1} \rangle = \langle v_{i+1}, z_{i+1} \rangle$$

if R is diagonal, and thus, in this case, (3.41) and (3.42) may be written instead as

$$\rho'_{i+1} = \rho'_i + \Delta\alpha_i \rho''_i - \langle x^s + w_{i+1} + \alpha_{i+1} v_{i+1}, z_{i+1} \rangle,$$

and

$$\rho''_{i+1} = \rho''_i - \langle v_{i+1}, z_{i+1} \rangle,$$

with no need to recur d_i in Step 4. Crucially, each of the recursions needed to maintain w_i , ρ_i , ρ'_i and ρ''_i only requires operations involving the sparse vector v_i .

To adapt the method for finding an approximate piecewise minimizer described in Section 2.1.2 to cope with regularization, the only significant issue is to consider how the regularization term $\rho(x)$ evolves as we move backwards or forwards along the search arc.

Plainly we have that

$$x_i := x^S + d_i = x^S + s_i^A + \alpha_i d_i^F = x_i^A + \alpha_i d_i^F,$$

where¹

$$x_i^A := x^S + s_i^A \quad \text{and} \quad d_i^F := I_{\mathcal{F}_i} d_{\mathcal{F}_i} \quad \text{with} \quad s_i^A := I_{\mathcal{A}_i} s_{\mathcal{A}_i},$$

and therefore

$$\rho(x_i) = \rho_i^C + \alpha_i \rho_i^L + \alpha_i^2 \rho_i^Q,$$

where

$$\rho_i^C := \frac{1}{2} \|x_i^A\|_R^2, \quad \rho_i^L := \langle x_i^A, d_i^F \rangle_R = \langle x^S, d_i^F \rangle_R \quad \text{and} \quad \rho_i^Q := \frac{1}{2} \|d_i^F\|_R^2.$$

It is then easy to simplify the recurrences described in the earlier section to deal with this. The generalisation of Algorithm 2.1 for the regularization case is then simply as follows.

Algorithm 3.2: Find an approximate backtracking arc minimizer x^C of ϕ

0. Initialization: The initial point $x^S \in \mathcal{X}$, search direction d , initial stepsize $\alpha_0 > 0$, reduction factor $\beta \in (0, 1)$ and decrease tolerance $\eta \in (0, \frac{1}{2})$ are given. Compute the residual $r^S = Ax^S - b$, the initial objective value $f(x^S) = \frac{1}{2} \|r^S\|_W^2$, the base fixed set \mathcal{A}^S from (2.21), the breakpoints α_j^B from (2.5) for all $j \in \mathbb{N}_n$, the end of the arc x^B from (2.22) and its direction $s = x^B - x^S$, the initial search point $x_0 = P_{\mathcal{X}}[x^S + \alpha_0 d]$, the active and free components at x_0 ,

$$\mathcal{A}_0 = \{j : \alpha_0 > \alpha_j^B\} \quad \text{and} \quad \mathcal{F}_0 = \{j \notin \mathcal{A}^S : \alpha_0 \leq \alpha_j^B\},$$

the corresponding values

$$x_0^A = x^S + s_0^A, \quad \text{with} \quad s_0^A = I_{\mathcal{A}_0} s_{\mathcal{A}_0} \quad \text{and} \quad d_0^F = I_{\mathcal{F}_0} d_{\mathcal{F}_0},$$

and residuals

$$r_0^A = r^S + p_0 \quad \text{and} \quad r_0^F = q_0$$

using the matrix-vector products

$$p_0 = A_{\mathcal{A}_0} s_{\mathcal{A}_0} \quad \text{and} \quad q_0 = A_{\mathcal{F}_0} d_{\mathcal{F}_0}.$$

Initialize

$$f_0^C = \frac{1}{2} \|r_0^A\|_W^2, \quad f_0^L = \langle r_0^A, r_0^F \rangle_W \quad \text{and} \quad f_0^Q = \frac{1}{2} \|r_0^F\|_W^2, \quad \text{and} \\ \rho_0^C = \frac{1}{2} \|x_0^A\|_R^2, \quad \rho_0^L = \langle x^S, d_0^F \rangle_R \quad \text{and} \quad \rho_0^Q = \frac{1}{2} \|d_0^F\|_R^2,$$

as well as

$$\gamma_0^A = \langle r^S, p_0 \rangle_W \quad \text{and} \quad \gamma_0^F = \langle r^S, q_0 \rangle_W, \quad \text{and} \\ \mu_0^A = \langle x^S, s_0^A \rangle_R \quad \text{and} \quad \mu_0^F = \langle x^S, d_0^F \rangle_R.$$

¹Here $s_{\mathcal{A}_i}$ and $d_{\mathcal{F}_i}$ are n -vectors, whose nonzero components are obtained from s and d via the index sets \mathcal{A}_i and \mathcal{F}_i , respectively.

Set $i = 0$.

1. Check for an approximate arc minimizer: Compute

$$\begin{aligned} f_i &= f_i^C + \alpha_i f_i^L + \alpha_i^2 f_i^Q, \\ \rho_i &= \rho_i^C + \alpha_i \rho_i^L + \alpha_i^2 \rho_i^Q, \\ \phi_i &= f_i + \sigma \rho_i \quad \text{and} \\ \gamma_i &= \gamma_i^A + \alpha_i \gamma_i^F + \sigma(\mu_i^A + \alpha_i \mu_i^F). \end{aligned}$$

If $\phi_i \leq \phi(x^S, \sigma) + \eta \gamma_i$, go to Step 4 if $i = 0$ but otherwise, i.e., if $i > 0$, set

$$\alpha^C = \alpha_i, \quad x^C = P_{\mathcal{X}}[x^S + \alpha^C d], \quad f(x^C) = f_i \quad \text{and} \quad \phi(x^C, \sigma) = \phi_i,$$

and stop.

2. Find the next set of indices that change status: Let $\alpha_{i+1} = \beta \alpha_i$, and compute

$$\mathcal{I}_{i+1} = \{j : \alpha_{i+1} < \alpha_j^B \leq \alpha_i\}$$

using the Heapsort algorithm.

3. Update the components of the objective and its slope: Compute

$$\begin{aligned} p_{i+1} &= A_{\mathcal{I}_{i+1}} s_{\mathcal{I}_{i+1}}, \quad y_{i+1} = W p_{i+1}, \\ q_{i+1} &= A_{\mathcal{I}_{i+1}} d_{\mathcal{I}_{i+1}} \quad \text{and} \quad z_{i+1} = R q_{i+1}, \\ s_{i+1} &= I_{\mathcal{I}_{i+1}} s_{\mathcal{I}_{i+1}}, \quad u_{i+1} = M s_{i+1}, \\ d_{i+1} &= I_{\mathcal{I}_{i+1}} d_{\mathcal{I}_{i+1}} \quad \text{and} \quad v_{i+1} = R s_{i+1}, \end{aligned}$$

update

$$\begin{aligned} f_{i+1}^C &= f_i^C - \langle y_{i+1}, r_i^A - \tfrac{1}{2} p_{i+1} \rangle, \\ f_{i+1}^L &= f_i^L + \langle z_{i+1}, r_i^A - \tfrac{1}{2} p_{i+1} \rangle - \langle y_{i+1}, r_i^F + \tfrac{1}{2} q_{i+1} \rangle, \\ f_{i+1}^Q &= f_i^Q + \langle z_{i+1}, r_i^F + \tfrac{1}{2} q_{i+1} \rangle, \\ \gamma_{i+1}^A &= \gamma_i^A - \langle y_{i+1}, r^S \rangle, \quad \gamma_{i+1}^F = \gamma_i^F + \langle z_{i+1}, r^S \rangle, \\ r_{i+1}^A &= r_i^A - p_{i+1} \quad \text{and} \quad r_{i+1}^F = r_i^F + q_{i+1}, \end{aligned}$$

as well as

$$\begin{aligned} \rho_{i+1}^C &= \rho_i^C - \langle u_{i+1}, x_i^A - \tfrac{1}{2} s_{i+1} \rangle, \\ \rho_{i+1}^L &= \rho_i^L + \langle v_{i+1}, x_i^A - \tfrac{1}{2} s_{i+1} \rangle - \langle u_{i+1}, d_i^F + \tfrac{1}{2} d_{i+1} \rangle, \\ \rho_{i+1}^Q &= \rho_i^Q + \langle v_{i+1}, d_i^F + \tfrac{1}{2} d_{i+1} \rangle, \\ \mu_{i+1}^A &= \mu_i^A - \langle u_{i+1}, x^S \rangle, \quad \mu_{i+1}^F = \mu_i^F + \langle v_{i+1}, x^S \rangle, \\ x_{i+1}^A &= x_i^A - s_{i+1} \quad \text{and} \quad d_{i+1}^F = d_i^F + d_{i+1}, \end{aligned}$$

increment i by 1 and return to Step 1.

4. Find the next set of indices that change status: Let $\alpha_{i+1} = \beta^{-1}\alpha_i$, and compute

$$\mathcal{J}_{i+1} = \{j : \alpha_i < \alpha_j^B \leq \alpha_{i+1}\}$$

using the Heapsort algorithm.

5. Update the components of the objective and its slope: Compute

$$\begin{aligned} p_{i+1} &= A_{\mathcal{J}_{i+1}} s_{\mathcal{J}_{i+1}}, \quad y_{i+1} = W p_{i+1}, \\ q_{i+1} &= A_{\mathcal{J}_{i+1}} d_{\mathcal{J}_{i+1}} \quad \text{and} \quad z_{i+1} = W q_{i+1}, \\ s_{i+1} &= I_{\mathcal{J}_{i+1}} s_{\mathcal{J}_{i+1}}, \quad u_{i+1} = M s_{i+1}, \\ d_{i+1} &= I_{\mathcal{J}_{i+1}} d_{\mathcal{J}_{i+1}} \quad \text{and} \quad v_{i+1} = R s_{i+1}, \end{aligned}$$

update

$$\begin{aligned} f_{i+1}^C &= f_i^C + \langle y_{i+1}, r_i^A + \tfrac{1}{2}p_{i+1} \rangle, \\ f_{i+1}^L &= f_i^L - \langle z_{i+1}, r_i^A + \tfrac{1}{2}p_{i+1} \rangle + \langle y_{i+1}, r_i^F - \tfrac{1}{2}q_{i+1} \rangle, \\ f_{i+1}^Q &= f_i^Q - \langle z_{i+1}, r_i^F - \tfrac{1}{2}q_{i+1} \rangle, \\ \gamma_{i+1}^A &= \gamma_i^A + \langle y_{i+1}, r^S \rangle, \quad \gamma_{i+1}^F = \gamma_i^F - \langle z_{i+1}, r^S \rangle, \\ r_{i+1}^A &= r_i^A + p_{i+1} \quad \text{and} \quad r_{i+1}^F = r_i^F - q_{i+1}, \end{aligned}$$

as well as

$$\begin{aligned} \rho_{i+1}^C &= \rho_i^C + \langle u_{i+1}, x_i^A + \tfrac{1}{2}d_{i+1} \rangle, \\ \rho_{i+1}^L &= \rho_i^L - \langle v_{i+1}, x_i^A + \tfrac{1}{2}d_{i+1} \rangle + \langle u_{i+1}, d_i^F - \tfrac{1}{2}d_{i+1} \rangle, \\ \rho_{i+1}^Q &= \rho_i^Q - \langle v_{i+1}, d_i^F - \tfrac{1}{2}d_{i+1} \rangle, \\ \mu_{i+1}^A &= \mu_i^A + \langle u_{i+1}, x^S \rangle, \quad \mu_{i+1}^F = \mu_i^F - \langle v_{i+1}, x^S \rangle, \\ x_{i+1}^A &= x_i^A + s_{i+1} \quad \text{and} \quad d_{i+1}^F = d_i^F - d_{i+1}. \end{aligned}$$

6. Check for an approximate extended arc minimizer: Compute

$$\begin{aligned} f_{i+1} &= f_{i+1}^C + \alpha_{i+1} f_{i+1}^L + \alpha_{i+1}^2 f_{i+1}^Q, \\ \rho_{i+1} &= \rho_{i+1}^C + \alpha_{i+1} \rho_{i+1}^L + \alpha_{i+1}^2 \rho_{i+1}^Q, \\ \phi_{i+1} &= f_{i+1} + \sigma \rho_{i+1} \quad \text{and} \\ \gamma_{i+1} &= \gamma_{i+1}^A + \alpha_{i+1} \gamma_{i+1}^F + \sigma(\mu_{i+1}^A + \alpha_{i+1} \mu_{i+1}^F). \end{aligned}$$

If $\phi_{i+1} > \phi(x^S, \sigma) + \eta \gamma_{i+1}$ or $\alpha_{i+1} \geq \max_j \alpha_j^B$, set

$$\alpha^C = \alpha_i, \quad x^C = P_{\mathcal{X}}[x^S + \alpha^C d], \quad f(x^C) = f_i \quad \text{and} \quad \phi(x^C, \sigma) = \phi_i,$$

and stop. Otherwise, increment i by 1 and return to Step 4.

3.1 Regularized linear least-squares minimization with fixed variables

Algorithm 3.3: The preconditioned conjugate-gradient regularized-least-squares method

Given x_0 , set $r_0 = Ax_0 - b$ and $g_0 = Z[A^T W r_0 + \sigma R x_0]$, and let $v_0 = Z[P^{-1}g_0]$ and $p_0 = -v_0$. For $k = 0, 1, \dots$ until convergence, perform the iteration

$$\begin{aligned} q_k &= Ap_k, \\ \alpha_k &= \langle g_k, v_k \rangle / (\langle q_k, W q_k \rangle + \sigma \langle p_k, R p_k \rangle), \\ x_{k+1} &= x_k + \alpha_k p_k, \\ r_{k+1} &= r_k + \alpha_k q_k, \\ g_{k+1} &= Z[A^T W r_{k+1} + \sigma R x_{k+1}], \\ v_{k+1} &= Z[P^{-1}g_{k+1}], \\ \beta_k &= \langle g_{k+1}, v_{k+1} \rangle / \langle g_k, v_k \rangle \text{ and} \\ p_{k+1} &= -v_{k+1} + \beta_k p_k. \end{aligned}$$

Notice that throughout only components $[p_k]_i$, $i \in \mathcal{F}$, can be nonzero, and this should be exploited when forming Ap_k and $\langle p_k, p_k \rangle_R$. The preconditioner needs to take account of the regularization term, and, at the very least, $P = \text{diag}(A^T W A) + \sigma \text{diag}(R)$ is appropriate.

Availability

The algorithms described have been implemented as the modern Fortran package `blls`, and the later is available as part of the GALAHAD library [4].

References

- [1] M. Bierlaire, Ph. L. Toint, and D. Tuytens. On iterative algorithms for linear least squares problems with bound constraints. *Linear Algebra and its Applications*, 143:111–143, 1991.
- [2] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. Testing a class of methods for solving minimization problems with simple bounds on the variables. *Mathematics of Computation*, 50:399–430, 1988.
- [3] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *Trust-Region Methods*. SIAM, Philadelphia, 2000.

- [4] N. I. M. Gould, D. Orban, and Ph. L. Toint. GALAHAD—a library of thread-safe fortran 90 packages for large-scale nonlinear optimization. *ACM Transactions on Mathematical Software*, 29(4):353–372, 2003.
- [5] N. I. M. Gould and D. P. Robinson. A dual gradient-projection method for large-scale strictly convex quadratic problems. *Computational Optimization and Applications*, 67(1):1–38, 2017.
- [6] N. I. M. Gould and J. A. Scott. The state-of-the-art of preconditioners for sparse linear least-squares problems. *ACM Transactions on Mathematical Software*, 43(4), 2017. Article 36.
- [7] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49:409–436, 1952.
- [8] J. J. Moré and G. Toraldo. On the solution of large quadratic programming problems with bound constraints. *SIAM Journal on Optimization*, 1(1):93–113, 1991.
- [9] C. C. Paige and M. A. Saunders. LSQR: an algorithm for sparse linear equations and sparse least squares. *ACM Transactions on Mathematical Software*, 8(1):43–71, 1982.
- [10] J. W. J. Williams. Algorithm 232, Heapsort. *Communications of the ACM*, 7:347–348, 1964.