A Constrained minimization algorithm

Before discussing our solution method for the linearly constrained ℓ_1-penalized least-squares problem, we briefly recall the homotopy/LARS method which manages to recover the unconstrained minimizer of the ℓ_1-penalized least-squares objective function

$$\bar{\mathbf{w}}(\tau) = \arg\min_{\mathbf{w}} \left[\|\boldsymbol{R}\mathbf{w} - \mathbf{y}\|_2^2 + \tau \|\mathbf{w}\|_1 \right]$$

for a whole range of values of the (positive) penalty parameter τ.

The variational equations describing the minimizer $\bar{\mathbf{w}}(\tau)$ are:

$$(\boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w}))_i = \frac{\tau}{2} \operatorname{sgn} w_i \qquad w_i \neq 0 \tag{1}$$

$$|(\boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w}))_i| \leq \frac{\tau}{2} \qquad w_i = 0. \tag{2}$$

The minimizer $\bar{\mathbf{w}}(\tau)$ is a continuous piecewise linear function of τ. We shall denote the breakpoints by $\tau_0 > \tau_1 > \dots$ and the corresponding minimizers by $\bar{\mathbf{w}}(\tau_0), \bar{\mathbf{w}}(\tau_1), \dots$ The breakpoints occur where a new component enters or leaves the support of $\bar{\mathbf{w}}(\tau)$. We will use $\mathbf{b}$ to denote the residual $\mathbf{b} = \boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w})$.

The homotopy/LARS method for solving these equations starts by considering the point $\mathbf{w} = 0$, which satisfies the equations (1,2) for all $\tau \geq \tau_0 \equiv 2 \max_i |(\boldsymbol{R}^\top \mathbf{y})_i|$. Hence $\bar{\mathbf{w}}(\tau \geq \tau_0) = 0$.

Given a breakpoint $\bar{\mathbf{w}}(\tau_n)$, it is possible to construct the next breakpoint $\bar{\mathbf{w}}(\tau_{n+1})$ by solving a small linear system. Let $J = \{i \text{ for which } |\mathbf{b}_i| = \tau_n/2\}$ (i.e. the set of maximal residual), $\boldsymbol{R}_J$ the submatrix consisting of the columns J of $\boldsymbol{R}$. We define the walking direction $\mathbf{u}$ by

$$\boldsymbol{R}_J^\top \boldsymbol{R}_J \, \mathbf{u}_J = \operatorname{sgn}(\mathbf{b}_J)$$

and $u_i = 0$ for $i \notin J$ ($\operatorname{sgn}(\mathbf{b}_J)$ denotes the vector $(\operatorname{sgn}(b_j)_{j \in J})$). In this way, a step $\mathbf{w} \to \mathbf{w} + \gamma \mathbf{u}$ results in a change in the residual $\mathbf{b} \to \mathbf{b} - \gamma \mathbf{v}$, where $v_j = \operatorname{sgn}(b_j)$ for $j \in J$. In other words, the maximal components of the residual decrease at the same rate. The step size $\gamma > 0$ is now determined to be the smallest number for which the absolute value of a component $|b_i - \gamma v_i|$ (with $i \notin J$) of the new residual becomes equal to $|b_j - \gamma v_j|$ for $j \in J$ (i.e. a new component joins the maximal residual set), or for which a nonzero component of $\mathbf{w}$ is turned into zero.

The new penalty parameter is then $\tau_{n+1} = \tau_n - 2\gamma$ (which is smaller than τ_n), and the corresponding minimizer is $\bar{\mathbf{w}}(\tau_{n+1}) = \bar{\mathbf{w}}(\tau_n) + \gamma \mathbf{u}$. By construction it is guaranteed to satisfy the variational equations (1,2).

The two main advantages of this method are thus that it is exact (in particular zero components are really zero) and that it yields the breakpoints (and hence the minimizers) for a whole range of values of the penalization parameters $\tau \geq \tau_{\text{stop}} \geq 0$. At each step, only a relatively small linear system has to be solved. If this procedure is carried through until the end, one finds $\lim_{\tau \to 0} \arg\min \|\boldsymbol{R}\mathbf{w} - \mathbf{y}\|_2^2 + \tau \|\mathbf{w}\|_1 = \arg\min_{\mathbf{w} \text{ s.t. } \boldsymbol{R}^\top \boldsymbol{R}\mathbf{w} = \boldsymbol{R}^\top \mathbf{y}} \|\mathbf{w}\|_1$.

For the constrained case, i.e. the minimization problem

$$\widetilde{\mathbf{w}}(\tau) = \arg\min_{\mathbf{w} \text{ s.t. } \boldsymbol{A}\mathbf{w} = \mathbf{a}} \left[\|\boldsymbol{R}\mathbf{w} - \mathbf{y}\|_2^2 + \tau \|\mathbf{w}\|_1 \right] \tag{3}$$

subject to the linear constraint $\boldsymbol{A}\mathbf{w} = \mathbf{a}$, we can devise a similar procedure. We assume, of course, that the constraint $\boldsymbol{A}\mathbf{w} = \mathbf{a}$ has a solution.

An approximation of the minimizer $\widetilde{\mathbf{w}}(\tau)$ can be obtained by applying the unconstrained procedure described above to the objective function

$$\widetilde{\bar{\mathbf{w}}}(\tau_\epsilon) = \arg\min_{\mathbf{w}} \left[\|\boldsymbol{A}\mathbf{w} - \mathbf{a}\|_2^2 + \epsilon \|\boldsymbol{R}\mathbf{w} - \mathbf{y}\|_2^2 + \tau_\epsilon \|\mathbf{w}\|_1 \right]. \tag{4}$$

For sufficiently small ϵ, this will give a good approximation of the constrained minimizer $\widetilde{\mathbf{w}}(\tau)$ corresponding to the penalty $\tau = \tau_\epsilon/\epsilon$ (after first going through a number of breakpoints for which $\boldsymbol{A}\mathbf{w} \neq \mathbf{a}$, not even approximately). However, this is clearly an approximate method (often very good) whereas the unconstrained procedure did not involve any approximation.

We solve this issue, and provide an exact method, by solving the minimization problem (4) up to the first order in ϵ. In this approach ϵ is a small *formal* positive parameter. Now the minimizer $\widetilde{\bar{\mathbf{w}}}(\tau_\epsilon)$ and τ_ϵ both depend on ϵ. We can write $\mathbf{w} = \mathbf{w}^{(0)} + \epsilon \mathbf{w}^{(1)} + \mathcal{O}(\epsilon^2)$ and $\tau_\epsilon = \tau^{(0)} + \tau^{(1)}\epsilon + \mathcal{O}(\epsilon^2)$. We again follow the procedure for the unconstrained method, but take care to use arithmetic (addition, multiplication, comparison, ...) up to first order in ϵ.

As before, one starts from $\mathbf{w} = 0$, corresponding to a large initial value of τ_ϵ, and follows the path of descending τ_ϵ. The strategy consists of satisfying the variational equations

$$\begin{aligned} \left(\boldsymbol{A}^\top(\mathbf{a} - \boldsymbol{A}\mathbf{w}) + \epsilon\boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w})\right)_i &= \frac{\tau_\epsilon}{2}\operatorname{sgn} w_i \qquad w_i \neq 0 \quad (5) \\ \left|\left(\boldsymbol{A}^\top(\mathbf{a} - \boldsymbol{A}\mathbf{w}) + \epsilon\boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w})\right)_i\right| &\leq \frac{\tau_\epsilon}{2} \qquad w_i = 0 \quad (6) \end{aligned}$$

at each breakpoint by carefully determining a walking direction $\mathbf{u} = \mathbf{u}^{(0)} + \mathbf{u}^{(1)}\epsilon + \mathcal{O}(\epsilon^2)$ and a step length $\gamma = \gamma^{(0)} + \gamma^{(1)}\epsilon + \mathcal{O}(\epsilon^2)$. Using $\mathbf{w} = \mathbf{w}^{(0)} + \epsilon\mathbf{w}^{(1)} + \mathcal{O}(\epsilon^2)$, we can rewrite equations (5) as

$$\begin{aligned} \left(\boldsymbol{A}^\top(\mathbf{a} - \boldsymbol{A}\mathbf{w}^{(0)})\right)_i &= \frac{\tau^{(0)}}{2}\operatorname{sgn} w_i \quad (7) \\ \left(-\boldsymbol{A}^\top\boldsymbol{A}\mathbf{w}^{(1)} + \boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w}^{(0)})\right)_i &= \frac{\tau^{(1)}}{2}\operatorname{sgn} w_i\ . \quad (8) \end{aligned}$$

From a known breakpoint $\mathbf{w}$ we can proceed to the following breakpoint by a step direction $\mathbf{u}$ and step size γ (both depending on ϵ). We again set

$$J = \arg\max_i \left|\left(\boldsymbol{A}^\top(\mathbf{a} - \boldsymbol{A}\mathbf{w}^{(0)}) + \epsilon(-\boldsymbol{A}^\top\boldsymbol{A}\mathbf{w}^{(1)} + \boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w}^{(0)}))\right)_i\right|.$$

As long as $\tau^{(0)} \neq 0$, the components J of $\mathbf{u}$ are determined by

$$\begin{pmatrix} \boldsymbol{R}_J^\top\boldsymbol{R}_J & \boldsymbol{A}_J^\top\boldsymbol{A}_J \\ \boldsymbol{A}_J^\top\boldsymbol{A}_J & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u}_J^{(0)} \\ \mathbf{u}_J^{(1)} \end{pmatrix} = \begin{pmatrix} 0 \\ \operatorname{sgn}(\mathbf{b}_J) \end{pmatrix} \quad (9)$$

and the other components of $\mathbf{u}$ remain zero. The step size γ is again determined as before, i.e. when a new component enters the maximal residual set, or when a component leaves the active set. The penalty parameter τ_ϵ decreases as before: $\tau_\epsilon \to \tau_\epsilon - 2\gamma$.

At some point in this procedure, τ_ϵ will become zero in zeroth order: $\tau_\epsilon = 0 + \tau^{(1)}\epsilon + \mathcal{O}(\epsilon^2)$. The corresponding minimizer (more precisely the zeroth-order part of this breakpoint) will satisfy the constraint $\boldsymbol{A}\mathbf{w} = \mathbf{a}$ and we will have found the first constrained minimizer $\widetilde{\mathbf{w}}$ of (3), corresponding to $\tau_0 = \tau^{(1)}$ (i.e. the first-order part of the parameter τ_ϵ of the ϵ-dependent problem at this breakpoint). In the unconstrained case, no such calculations were necessary as the starting point was always equal to 0. Similarly to the unconstrained case, we have that $\widetilde{\mathbf{w}}(\tau > \tau_0) = \widetilde{\mathbf{w}}(\tau_0)$.

In principle, one could continue the ϵ-dependent algorithm, but now that the first breakpoint of $\widetilde{\mathbf{w}}(\tau)$ is determined, it is more advantageous to continue the descent of τ by introducing Lagrange multipliers $\boldsymbol{\lambda}$ for the problem (3):

$$\widetilde{\mathbf{w}}(\tau) = \underset{\boldsymbol{\lambda},\ \mathbf{w}\ \text{s.t.}\ \boldsymbol{A}\mathbf{w}=\mathbf{a}}{\arg\min} \left[\|\boldsymbol{R}\mathbf{w} - \mathbf{y}\|_2^2 + \tau\|\mathbf{w}\|_1 + 2\boldsymbol{\lambda}^\top(\boldsymbol{A}\mathbf{w} - \mathbf{a})\right].$$

This minimization problem (analogous to (3)) amounts to solving the equations:

$$\begin{aligned} (\boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w}) + \boldsymbol{A}^\top\boldsymbol{\lambda})_i &= \frac{\tau}{2}\operatorname{sgn} w_i \qquad w_i \neq 0 \quad (10) \\ |(\boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w}) + \boldsymbol{A}^\top\boldsymbol{\lambda})_i| &\leq \frac{\tau}{2} \qquad w_i = 0 \quad (11) \\ \boldsymbol{A}\mathbf{w} &= \mathbf{a}. \quad (12) \end{aligned}$$

Equation (10) is the equivalent of equation (8) whereas equation (12) replaces equation (7). We now already have τ_0, $\widetilde{\mathbf{w}}(\tau \geq \tau^{(0)})$ and the initial Lagrange multipliers $\boldsymbol{\lambda} = -\boldsymbol{A}\widetilde{\widetilde{\mathbf{w}}}^{(1)}$ (from the first-order part of the last step of the ϵ-dependent problem).

To proceed from one breakpoint to the next ($\mathbf{w} \to \mathbf{w} + \gamma\mathbf{u}$, and $\boldsymbol{\lambda} \to \boldsymbol{\lambda} + \gamma\mathbf{s}$ as the multipliers also change), we again need to solve a linear system:

$$\begin{pmatrix} \boldsymbol{R}_J^\top\boldsymbol{R}_J & \boldsymbol{A}_J^\top \\ \boldsymbol{A}_J & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u}_J \\ \mathbf{s} \end{pmatrix} = \begin{pmatrix} \operatorname{sgn}(\tilde{\mathbf{b}}_J) \\ 0 \end{pmatrix} \quad (13)$$

with $\tilde{\mathbf{b}} = \boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w}) + \boldsymbol{A}^\top\boldsymbol{\lambda}$. This will guarantee that $\mathbf{w} \to \mathbf{w} + \gamma\mathbf{u}$ and $\boldsymbol{\lambda} \to \boldsymbol{\lambda} + \gamma\mathbf{s}$ still satisfy the constraint (12) and the variational equations (10,11). The step size γ is determined by the same rule as before: stop when a new component enters the set $J = \arg\max_i |(\boldsymbol{R}^\top(\mathbf{y} - \boldsymbol{R}\mathbf{w}) + \boldsymbol{A}^\top\boldsymbol{\lambda})_i|$ or when a nonzero component of $\mathbf{w}$ is set to zero. Notice the differences and similarities between the linear systems (13) and (9).

At each breakpoint, this algorithm provides the penalty τ_n, the corresponding minimizer $\widetilde{\mathbf{w}}(\tau_n)$ and the Lagrange multipliers $\boldsymbol{\lambda}_n$. Unlike for the unconstrained case, it is now possible that $\widetilde{\mathbf{w}}(\tau)$ remains constant between two breakpoints (i.e. only the Lagrange multipliers $\boldsymbol{\lambda}$ change).

One simplifying assumption (not solved in the homotopy/LARS algorithm) was made in the above description of the algorithm: if the set of maximal residual and the support set differ by more than one component, one should carefully select the correct new components to enter the support. This can be done by using the variational equations, and our implementation handles this case.

One could argue that the starting point (i.e. the first breakpoint) for the constrained minimization problem is simply given by $\widetilde{\mathbf{w}}(\tau_0) = \underset{\mathbf{w}\ \text{s.t.}\ \boldsymbol{A}\mathbf{w}=\mathbf{a}}{\arg\min} \|\mathbf{w}\|_1$, which could be calculated by letting the unconstrained solution procedure run its course: $\widetilde{\mathbf{w}}(\tau_0) = \lim_{\sigma\to 0} \arg\min_{\mathbf{w}} \left[\|\boldsymbol{A}\mathbf{w} - \mathbf{a}\|_2^2 + \sigma\|\mathbf{w}\|_1\right]$. Generically (i.e. excluding special cases), this is correct. However, the problem is that sometimes the minimizer $\operatorname{argmin}_{\mathbf{w}\ \text{s.t.}\ \boldsymbol{A}\mathbf{w}=\mathbf{a}} \|\mathbf{w}\|_1$ is not unique. In that case, the starting point for the constrained minimizer is not solely determined by $\boldsymbol{A}$ and $\mathbf{a}$ but also by $\boldsymbol{R}$ and $\mathbf{y}$. In this case, the ϵ-dependent algorithm still chooses the correct starting point from the set $\underset{\mathbf{w}\ \text{s.t.}\ \boldsymbol{A}\mathbf{w}=\mathbf{a}}{\arg\min} \|\mathbf{w}\|_1$. This is important to mention because the special constraint $\sum w_i = 1$ used in this paper, gives rise to such cases.

Our algorithm is well-suited for the portfolio problems discussed in this paper. The size of the matrix, the number of constraints (just two) and, more importantly, the number of nonzero weights in the portfolios are such that a minimization run (i.e. finding the minimizer for a whole range of penalty parameters) can be done in a fraction of a second on a standard desktop.

We calculated the portfolio examples in this paper using both the formal ϵ approach (in Mathematica) and the approximate small ϵ approach (in Matlab). The outcomes were always consistent.