

Coordinate descent

(Coordinate-wise minimization)

Barnabas Póczos & Ryan Tibshirani
Convex Optimization 10-725/36-725

Coordinate descent


We've seen some pretty sophisticated methods

Our focus today is a very simple technique that can be surprisingly efficient and scalable: **coordinate descent**, i.e., coordinate-wise minimization

Q: Given convex, differentiable $f : \mathbb{R}^n \rightarrow \mathbb{R}$, if we are at a point x such that $f(x)$ is minimized along each coordinate axis, *have we found a global minimizer?*

I.e., does $f(x + \delta \cdot e_i) \geq f(x)$ for all $\delta, i \implies f(x) = \min_z f(z)$?



(Here $e_i = (0, \dots, 1, \dots, 0) \in \mathbb{R}^n$, the i th standard basis vector)



A: Yes! Proof:



$$\nabla f(x) = \left(\frac{\partial f}{\partial x_1}(x), \dots, \frac{\partial f}{\partial x_n}(x) \right) = 0$$

Q: Same question, but for f convex (not differentiable) ... ?



A: No! Look at the above counterexample

Q: Same question again, but now $f(x) = g(x) + \sum_{i=1}^n h_i(x_i)$, with g convex, differentiable and each h_i convex ... ? (Non-smooth part here called **separable**)



A: Yes! Proof: for any y ,

$$\begin{aligned}
 f(y) - f(x) &\geq \nabla g(x)^T (y - x) + \sum_{i=1}^n [h_i(y_i) - h_i(x_i)] \\
 &= \sum_{i=1}^n \underbrace{[\nabla_i g(x)(y_i - x_i) + h_i(y_i) - h_i(x_i)]}_{\geq 0} \geq 0
 \end{aligned}$$

Coordinate descent

This suggests that for $f(x) = g(x) + \sum_{i=1}^n h_i(x_i)$ (with g convex, differentiable and each h_i convex) we can use **coordinate descent** to find a minimizer: start with some initial guess $x^{(0)}$, and repeat

$$x_1^{(k)} \in \operatorname{argmin}_{x_1} f(x_1, x_2^{(k-1)}, x_3^{(k-1)}, \dots, x_n^{(k-1)})$$

$$x_2^{(k)} \in \operatorname{argmin}_{x_2} f(x_1^{(k)}, x_2, x_3^{(k-1)}, \dots, x_n^{(k-1)})$$

$$x_3^{(k)} \in \operatorname{argmin}_{x_3} f(x_1^{(k)}, x_2^{(k)}, x_3, \dots, x_n^{(k-1)})$$

...

$$x_n^{(k)} \in \operatorname{argmin}_{x_n} f(x_1^{(k)}, x_2^{(k)}, x_3^{(k)}, \dots, x_n)$$

for $k = 1, 2, 3, \dots$ (note: after we solve for $x_i^{(k)}$, we use its new value from then on!)

Seminal work of Tseng (2001) proves that for such f (provided f is continuous on compact set $\{x : f(x) \leq f(x^{(0)})\}$ and f attains its minimum), any limit point of $x^{(k)}$, $k = 1, 2, 3, \dots$ is a minimizer of f^1

Notes:

- Order of cycle through coordinates is arbitrary, can use any permutation of $\{1, 2, \dots, n\}$
- Can everywhere replace individual coordinates with blocks of coordinates
- “One-at-a-time” update scheme is critical, and “all-at-once” scheme **does not** necessarily converge

¹Using real analysis, we know that $x^{(k)}$ has subsequence converging to x^* (Bolzano-Weierstrass), and $f(x^{(k)})$ converges to f^* (monotone convergence)

Linear regression

Consider linear regression

$$\min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2$$

where $y \in \mathbb{R}^n$, and $X \in \mathbb{R}^{n \times p}$ with columns X_1, \dots, X_p

Minimizing over β_i , with all β_j , $j \neq i$ fixed:


$$0 = \nabla_i f(\beta) = X_i^T (X\beta - y) = X_i^T (X_i\beta_i + X_{-i}\beta_{-i} - y)$$

i.e., we take

$$\beta_i = \frac{X_i^T (y - X_{-i}\beta_{-i})}{X_i^T X_i}$$

Coordinate descent repeats this update for $i = 1, 2, \dots, p, 1, 2, \dots$


Coordinate descent vs gradient descent for linear regression: 100 instances ($n = 100, p = 20$)



Is it fair to compare 1 cycle of coordinate descent to 1 iteration of gradient descent? Yes, if we're clever:

$$\beta_i \leftarrow \frac{X_i^T (y - X_{-i} \beta_{-i})}{X_i^T X_i} = \frac{X_i^T r}{\|X_i\|_2^2} + \beta_i$$

where $r = y - X\beta$. Therefore each coordinate update takes $O(n)$ operations — $O(n)$ to update r , and $O(n)$ to compute $X_i^T r$ — and one cycle requires $O(np)$ operations, just like gradient descent



Same example, but now with accelerated gradient descent for comparison

Is this contradicting the optimality of accelerated gradient descent? I.e., is coordinate descent a first-order method?

No. It uses much more than first-order information

Lasso regression

Now consider the lasso problem

$$\min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 + \lambda \|\beta\|_1$$

Note that the non-smooth part is separable: $\|\beta\|_1 = \sum_{i=1}^p |\beta_i|$

Minimizing over β_i , with β_j , $j \neq i$ fixed:

$$0 = X_i^T X_i \beta_i + X_i^T (X_{-i} \beta_{-i} - y) + \lambda s_i$$

where $s_i \in \partial |\beta_i|$. Solution is simply given by soft-thresholding

$$\beta_i = S_{\lambda / \|X_i\|_2} \left(\frac{X_i^T (y - X_{-i} \beta_{-i})}{X_i^T X_i} \right)$$

Repeat this for $i = 1, 2, \dots, p, 1, 2, \dots$

Box-constrained regression

Consider box-constrained linear regression

$$\min_{\beta \in \mathbb{R}^p} \frac{1}{2} \|y - X\beta\|_2^2 \quad \text{subject to} \quad \|\beta\|_\infty \leq s$$

Note this fits our framework, as $1\{\|\beta\|_\infty \leq s\} = \sum_{i=1}^n 1\{|\beta_i| \leq s\}$

Minimizing over β_i with all β_j , $j \neq i$ fixed: same basic steps give

$$\beta_i = T_s \left(\frac{X_i^T (y - X_{-i} \beta_{-i})}{X_i^T X_i} \right)$$

where T_s is the truncating operator:

$$T_s(u) = \begin{cases} s & \text{if } u > s \\ u & \text{if } -s \leq u \leq s \\ -s & \text{if } u < -s \end{cases}$$

Support vector machines

A coordinate descent strategy can be applied to the SVM dual:

$$\min_{\alpha \in \mathbb{R}^n} \frac{1}{2} \alpha^T K \alpha - 1^T \alpha \quad \text{subject to} \quad y^T \alpha = 0, \quad 0 \leq \alpha \leq C \mathbf{1}$$

Sequential minimal optimization or SMO (Platt 1998) is basically blockwise coordinate descent in blocks of 2. Instead of cycling, it chooses the next block greedily

Recall the complementary slackness conditions

$$\alpha_i \cdot [(Av)_i - y_i d - (1 - s_i)] = 0, \quad i = 1, \dots, n \quad (1)$$


$$(C - \alpha_i) \cdot s_i = 0, \quad i = 1, \dots, n \quad (2)$$

where v, d, s are the primal coefficients, intercept, and slacks, with $v = A^T \alpha$, d computed from (1) using any i such that $0 < \alpha_i < C$, and s computed from (1), (2)

SMO repeats the following two steps:

- Choose α_i, α_j that do not satisfy complementary slackness
- Minimize over α_i, α_j exactly, keeping all other variables fixed

Second step uses equality constraint, reduces to minimizing univariate quadratic over an interval (From Platt 1998)



First step uses heuristics to choose α_i, α_j greedily

Note this does not meet separability assumptions for convergence from Tseng (2001), and a different treatment is required

Coordinate descent in statistics and ML

History in statistics:

- Idea appeared in Fu (1998), and again in Daubechies et al. (2004), but was inexplicably ignored
- Three papers around 2007, and Friedman et al. (2007) really sparked interest in statistics and ML community

Why is it used?

- Very simple and easy to implement
- Careful implementations can attain state-of-the-art
- Scalable, e.g., don't need to keep data in memory

Some examples: lasso regression, SVMs, lasso GLMs, group lasso, fused lasso², trend filtering², graphical lasso², and even regression with nonconvex penalties

²Coordinate descent here is typically applied to the dual problem

Pathwise coordinate descent for lasso

Here is the basic outline for pathwise coordinate descent for lasso, from Friedman et al. (2007), Friedman et al. (2009)

Outer loop (**pathwise** strategy):

- Compute the solution over a sequence $\lambda_1 > \lambda_2 > \dots > \lambda_r$ of tuning parameter values
- For tuning parameter value λ_k , initialize coordinate descent algorithm at the computed solution for λ_{k+1} (warm start)

Inner loop (**active set** strategy):

- Perform one coordinate cycle (or small number of cycles), and record active set S of coefficients that are nonzero
- Cycle over coefficients in S until convergence
- Check KKT conditions over all coefficients; if not all satisfied, add offending coefficients to S , go back one step

Even when the solution is only desired at one value of λ , pathwise strategy ($\lambda_1 > \lambda_2 > \dots > \lambda_r = \lambda$) is typically much more efficient than directly performing coordinate descent at λ

Active set strategy takes advantage of sparsity; e.g., for very large problems, coordinate descent for lasso is much faster than it is for ridge regression

With these strategies in place (and a few more tricks), coordinate descent can be competitive with fastest algorithms for ℓ_1 penalized minimization problems

Freely available via [glmnet](#) package in MATLAB or R (Friedman et al. 2009)

Convergence rates?

Global convergence rates for coordinate descent have not yet been established as they have for first-order methods

Recently Saha et al. (2013) consider minimizing

$$f(x) = g(x) + \lambda \|x\|_1$$

and assume that

- g convex, ∇g Lipschitz with constant $L > 0$, and $I - \nabla g/L$ monotone increasing in each component
- there is z such that $z \geq S_\lambda(z - \nabla g(z))$ or $z \leq S_\lambda(z - \nabla g(z))$ (component-wise)

They show that for coordinate descent starting at $x^{(0)} = z$, and generalized gradient descent starting at $y^{(0)} = z$ (step size $1/L$),

$$f(x^{(k)}) - f(x^*) \leq f(y^{(k)}) - f(x^*) \leq \frac{L \|x^{(0)} - x^*\|_2^2}{2k}$$

Graphical lasso

Consider a data matrix $X \in \mathbb{R}^{n \times p}$, whose rows $x_1, \dots, x_n \in \mathbb{R}^p$ are independent observations from $N(0, \Sigma)$, with unknown covariance matrix Σ . Want to estimate Σ

For $Z \sim N(0, \Sigma)$, normality theory tells us

$$\Sigma_{ij}^{-1} = 0 \iff Z_i, Z_j \text{ conditionally independent given } Z_\ell, \ell \neq i, j$$


If we believe that many components are conditionally independent given others (often reasonable for large p) then we want a sparse estimate of Σ^{-1} . Can get this by solving **graphical lasso** problem (Banerjee et al. 2007, Friedman et al. 2007):

$$\min_{\Theta \in \mathbb{R}^{p \times p}, \Theta \succ 0} -\log \det \Theta + \text{tr}(S\Theta) + \lambda \|\Theta\|_1$$


Minimizer $\hat{\Theta}$ is an estimate for Σ^{-1} . (Note here $S = X^T X/n$ is the empirical covariance matrix, and $\|\Theta\|_1 = \sum_{i,j=1}^p |\Theta_{ij}|$)

Example from Friedman et al. (2007), cell-signaling network:

Believed network




Graphical lasso estimates




Example from Liu et al. (2010), hub graph simulation:

True graph



Graphical lasso estimate



Graphical lasso KKT conditions (stationarity):

$$-\Theta^{-1} + S + \lambda\Gamma = 0$$

where $\Gamma_{ij} \in \partial|\Theta_{ij}|$. Let $W = \Theta^{-1}$; we will solve in terms of W . Note $W_{ii} = S_{ii} + \lambda$, because $\Theta_{ii} > 0$ at solution. Now partition:

$$\begin{array}{cccc} W = & \Theta = & S = & \Gamma = \\ \left[\begin{array}{cc} W_{11} & w_{12} \\ w_{21} & w_{22} \end{array} \right] & \left[\begin{array}{cc} \Theta_{11} & \theta_{12} \\ \theta_{21} & \theta_{22} \end{array} \right] & \left[\begin{array}{cc} S_{11} & s_{12} \\ s_{21} & s_{22} \end{array} \right] & \left[\begin{array}{cc} \Gamma_{11} & \gamma_{12} \\ \gamma_{21} & \gamma_{22} \end{array} \right] \end{array}$$

where $W_{11} \in \mathbb{R}^{(p-1) \times (p-1)}$, $w_{12} \in \mathbb{R}^{(p-1) \times 1}$, and $w_{21} \in \mathbb{R}^{1 \times (p-1)}$, $w_{22} \in \mathbb{R}$; same with others

Coordinate descent strategy: solve for w_{12} , the last column of W (note w_{22} is known), with all other columns fixed; then solve for second-to-last column, etc., and cycle around until convergence. (Solve for Θ along the way, so we don't have to invert W to get Θ)

Now consider (1, 2)-block of KKT conditions:

$$-w_{12} + s_{12} + \lambda\gamma_{12} = 0$$

Because $\begin{bmatrix} W_{11} & w_{12} \\ w_{21} & w_{22} \end{bmatrix} \begin{bmatrix} \Theta_{11} & \theta_{12} \\ \theta_{21} & \theta_{22} \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & 1 \end{bmatrix}$, we know that

$w_{12} = -W_{11}\theta_{12}/\theta_{22}$. Substituting this into the above,

$$W_{11} \frac{\theta_{12}}{\theta_{22}} + s_{12} + \lambda\gamma_{12} = 0$$

Letting $\beta = \theta_{12}/\theta_{22}$ and noting that $\theta_{22} > 0$ at solution, this is

$$W_{11}\beta + s_{12} + \lambda\rho = 0$$

where $\rho \in \partial\|\beta\|_1$. What does this condition look like?

These are exactly the KKT conditions for

$$\min_{\beta \in \mathbb{R}^{p-1}} \beta^T W_{11} \beta + s_{12}^T \beta + \lambda \|\beta\|_1$$

which is (basically) a lasso problem and can be itself solved quickly via coordinate descent

From β we get $w_{12} = -W_{11}\beta$, and set $w_{21} = w_{12}^T$. Then θ_{12}, θ_{22}

are obtained from $\begin{bmatrix} W_{11} & w_{12} \\ w_{21} & w_{22} \end{bmatrix} \begin{bmatrix} \Theta_{11} & \theta_{12} \\ \theta_{21} & \theta_{22} \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & 1 \end{bmatrix}$, and we set $\theta_{21} = \theta_{12}^T$


The next step moves on to a different column of W , and so on; hence we have reduced the graphical lasso problem to a repeated sequence of lasso problems

This coordinate descent approach for the graphical lasso, usually called **glasso** algorithm (Friedman et al. 2007) is very efficient and scales well

Meanwhile, people have noticed that using glasso algorithm, it can happen that the objective function doesn't decrease monotonically across iterations — is this a bug?

No! The glasso algorithm makes a variable transformation and solves in terms of coordinate blocks of W ; note that these are **not coordinate blocks** of original variable Θ , so strictly speaking it is not a coordinate descent algorithm

However, it can be shown that glasso is doing coordinate ascent on the dual problem (Mazumder et al. 2011)



Screening rules for graphical lasso

Graphical lasso computations can be significantly accelerated by using a clever screening rule (this is analogous to the SAFE rules for the lasso)

Mazumder et al. (2011), Witten et al. (2011) examine the KKT conditions:

$$-\Theta^{-1} + S + \lambda\Gamma = 0$$

and conclude that Θ is block diagonal over variables C_1, C_2 if and only if $|S_{ij}| \leq \lambda$ for all $i \in C_1, j \in C_2$. Why?

- If Θ is block diagonal, then so is Θ^{-1} , and thus $|S_{ij}| \leq \lambda$ for $i \in C_1, j \in C_2$
- If $|S_{ij}| \leq \lambda$ for $i \in C_1, j \in C_2$, then the KKT conditions are satisfied with Θ^{-1} block diagonal, so Θ is block diagonal

Exact same idea extends to multiple blocks. Hence group structure in graphical lasso solution is just given by **covariance thresholding**

References

Early coordinate descent references in statistics and ML:

- I. Daubechies and M. Defrise and C. De Mol (2004), “An iterative thresholding algorithm for linear inverse problems with a sparsity constraint”
- J. Friedman and T. Hastie and H. Hoefling and R. Tibshirani (2007), “Pathwise coordinate optimization”
- W. Fu (1998), “Penalized regressions: the bridge versus the lasso”
- T. Wu and K. Lange (2008), “Coordinate descent algorithms for lasso penalized regression”
- A. van der Kooij (2007), “Prediction accuracy and stability of regression with optimal scaling transformations”

Applications of coordinate descent:

- O. Banerjee and L. Ghaoui and A. d'Aspremont (2007), "Model selection through sparse maximum likelihood estimation"
- J. Friedman and T. Hastie and R. Tibshirani (2007), "Sparse inverse covariance estimation with the graphical lasso"
- J. Friedman and T. Hastie and R. Tibshirani (2009), "Regularization paths for generalized linear models via coordinate descent"
- J. Platt (1998), "Sequential minimal optimization: a fast algorithm for training support vector machines"

Theory for coordinate descent:

- R. Mazumder and J. Friedman and T. Hastie (2011), "SparseNet: coordinate descent with non-convex penalties"
- A. Saha and A. Tewari (2013), "On the nonasymptotic convergence of cyclic coordinate descent methods"
- P. Tseng (2001), "Convergence of a block coordinate descent method for nondifferentiable minimization"

More graphical lasso references:

- H. Liu and K. Roeder and L. Wasserman (2010), “Stability approach to regularization selection (StARS) for high dimensional graphical models”
- R. Mazumder and T. Hastie (2011), “The graphical lasso: new insights and alternatives”
- R. Mazumder and T. Hastie (2011), “Exact covariance thresholding into connected components for large-scale graphical lasso”
- D. Witten and J. Friedman and N. Simon (2011), “New insights and faster computations for the graphical lasso”