A SKETCH OF THE CONVERGENCE OF STEEPEST DESCENT METHOD FOR QUADRATIC OBJECTIVES

David F. Gleich

February 6, 2023

We'll study the steepest descent method, also known as the gradient descent method, on a simple quadratic objective with exact line search. The point here is to show that, even in a simple case, this method converges slowly. (In this class, that means linear convergence.) Whether or not this slow convergence affects your problem is a judgement call that you'll have to make for yourself.

This material comes from Nocedal & Wright, page 42; Griva, Sofer & Nash, page 405; and some old notes from Juan Meza (former location: `http://hpcrd.lbl.gov/~meza/papers/Steepest%20Descent.pdf`) which seem to be in a Wiley paper `http://dx.doi.org/10.1002/wics.117`.

1 THE PROBLEM

We'll consider the optimization problem

$$\text{minimize} \quad \tfrac{1}{2}\mathbf{x}^T \boldsymbol{Q} \mathbf{x}$$

where $\boldsymbol{Q}$ is symmetric, positive definite.

The solution of this problem is $\mathbf{x} = 0$. We'll show that it takes us a while to find this solution using steepest descent!

This is a strongly convex problem with a unique solution – just about as easy as things get.

2 THE METHOD

Steepest descent begins with some prescribed point $\mathbf{x}_0$. At each step, it considers a linear approximation of $f(\mathbf{x})$ in the direction of the negative gradient. Formally,

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha_k \mathbf{g}_k$$

where $\mathbf{g}_k = \mathbf{g}(\mathbf{x}_k)$ is the gradient evaluated at $\mathbf{x}_k$.

The key decision in the method is how to choose α_k, what is called the step-length. An idealized choice of α_k is as the global solution of:

$$\begin{array}{ll} \underset{\alpha}{\text{minimize}} & f(\mathbf{x}_k - \alpha \mathbf{g}_k) \\ \text{subject to} & \alpha > 0 \end{array} .$$

The negative gradient search direction is optimal for the model objective, $f(\mathbf{x}) \approx f(\mathbf{x}_k + \mathbf{p}) \approx f(\mathbf{x}_k) + \mathbf{p}^T \mathbf{g}(\mathbf{x}_k)$, in the sense that $-\mathbf{g}(\mathbf{x}_k)$ is the solution of

$$\text{minimize} \quad \frac{\mathbf{p}^T \mathbf{g}_k}{\|\mathbf{p}\| \|\mathbf{g}_k\|} .$$

For reasons that we'll see soon, this is called *exact line search*. Usually, performing an exact line search is impossible, but for quadratic objectives of the form:

$$f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \boldsymbol{Q}\mathbf{x} - \beta \mathbf{x}^T \mathbf{c}$$

we can derived a closed form solution for α.

3 EXACT LINE SEARCH FOR QUADRATICS

Let's do so for the simple objective $f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \boldsymbol{Q}\mathbf{x}$. First, we need the gradient. For this function $\mathbf{g}(\mathbf{x}) = \boldsymbol{Q}\mathbf{x}$ if $\boldsymbol{Q}$ is symmetric, and $(\frac{1}{2}\boldsymbol{Q} + \frac{1}{2}\boldsymbol{Q}^T)\mathbf{x}$ if $\boldsymbol{Q}$ is non-symmetric. We'll only consider the symmetric case where $\boldsymbol{Q}$ is also positive definite.

At a point $\mathbf{x}_k$, then, our goal is to pick α_k to minimize

$$\ell(\alpha) = f(\mathbf{x}_k - \alpha \mathbf{g}_k) = \tfrac{1}{2}(\mathbf{x}_k - \alpha \mathbf{g}_k)^T \boldsymbol{Q}(\mathbf{x}_k - \alpha \mathbf{g}_k) = \tfrac{1}{2}\mathbf{x}_k^T \boldsymbol{Q}\mathbf{x}_k - \alpha \mathbf{g}_k^T \boldsymbol{Q}\mathbf{x}_k + \tfrac{1}{2}\alpha^2 \mathbf{g}_k^T \boldsymbol{Q}\mathbf{g}_k .$$

The derivative with respect to α is:

$$\ell'(\alpha) = \alpha \mathbf{g}_k^T \boldsymbol{Q}\mathbf{g}_k - \mathbf{g}_k^T \boldsymbol{Q}\mathbf{x}_k = \alpha \mathbf{g}_k^T \boldsymbol{Q}\mathbf{g}_k - \mathbf{g}_k^T \mathbf{g}_k .$$

The only stationary point is where $\ell'(\alpha) = 0$, or

$$\alpha = \frac{\|\mathbf{g}_k\|^2}{\mathbf{g}_k^T \mathbf{Q} \mathbf{g}_k}.$$

Because of the norms, this ratio must be positive (unless $\mathbf{g}_k = 0$, in which case we are already at a stationary point). Thus, the solution is feasible and hence, optimal.

4 DECREASE IN THE FUNCTION

Our goal is to understand the convergence properties of the steepest descent method. Now that we have fully specified the method by deriving the form of α at each step, we can study the decrease in the objective function from $f(\mathbf{x}_k)$ to $f(\mathbf{x}_{k+1})$. Substituting $\alpha_k = \frac{\|\mathbf{g}_k\|^2}{\mathbf{g}_k^T \mathbf{Q} \mathbf{g}_k}$ into $f(\mathbf{x}_k - \alpha \mathbf{g}_k)$ we have:

$$f(\mathbf{x}_{k+1}) = f(\mathbf{x}_k) - \alpha \|\mathbf{g}\|^2 + \tfrac{1}{2}\alpha^2 \mathbf{g}_k^T \mathbf{Q} \mathbf{g}_k = f(\mathbf{x}_k) - \tfrac{1}{2}\|\mathbf{g}_k\|^4 / (\mathbf{g}_k^T \mathbf{Q} \mathbf{g}_k).$$

From this form, it seems like we might have algebraic convergence! However, note that

$$f(\mathbf{x}_k) = \tfrac{1}{2}\mathbf{x}_k^T \mathbf{Q} \mathbf{x}_k = \tfrac{1}{2}\mathbf{g}_k^T \mathbf{Q}^{-1} \mathbf{g}_k.$$

Thus, we can rewrite the decrease as follows:

$$f(\mathbf{x}_{k+1}) = f(\mathbf{x}_k)\left[1 - \frac{\|\mathbf{g}_k\|^4}{(\mathbf{g}_k^T \mathbf{Q} \mathbf{g}_k)(\mathbf{g}_k^T \mathbf{Q}^{-1} \mathbf{g}_k)}\right].$$

If we can bound $\frac{\|\mathbf{g}_k\|^4}{(\mathbf{g}_k^T \mathbf{Q} \mathbf{g}_k)(\mathbf{g}_k^T \mathbf{Q}^{-1} \mathbf{g}_k)}$ by a constant, then we will have proved linear convergence.

5 THE KANTOROVICH INEQUALITY

There is indeed such a bound, known as the Kantorovich inequality.

LEMMA 1 *Let $\boldsymbol{A}$ be a symmetric, positive definite matrix. Let $\mathbf{x}$ be a vector with $\mathbf{x}^T\mathbf{x} = 1$. Then*

$$(\mathbf{x}^T \boldsymbol{A} \mathbf{x})(\mathbf{x}^T \boldsymbol{A}^{-1} \mathbf{x}) \le \frac{(\lambda_1 + \lambda_n)^2}{4\lambda_1\lambda_n}$$

where λ_1 and λ_n are the smallest and largest eigenvalues of $\boldsymbol{A}$.

This proof due to P. Henrici, Mathematical Notes, 1961.

Proof First, we transform the problem into eigenvalue space with $\boldsymbol{A} = \boldsymbol{V}\boldsymbol{\Lambda}\boldsymbol{V}^T$. This results in

$$(\mathbf{x}^T \boldsymbol{A} \mathbf{x})(\mathbf{x}^T \boldsymbol{A}^{-1} \mathbf{x}) = (\mathbf{y}^T \boldsymbol{\Lambda} \mathbf{y})(\mathbf{y}^T \boldsymbol{\Lambda}^{-1} \mathbf{y}) = \left(\sum_{i=1}^{n} y_i^2 \lambda_i\right)\left(\sum_{i=1}^{n} y_i^2 \lambda_i^{-1}\right).$$

Note that $\sum_i y_i^2 = 1$. The key idea is to write $\lambda_i = p_i\lambda_1 + q_i\lambda_n$ and $\lambda_i^{-1} = p_i\lambda_1^{-1} + q_i\lambda_n^{-1}$. Thus, $p_i \ge 0$ and $q_i \ge 0$, and

$$\lambda_i\lambda_i^{-1} = 1 = (p_i + q_i)^2 + p_i q_i(\lambda_1 - \lambda_n)^2/(\lambda_1\lambda_n).$$

Let $p = \sum_i y_i^2 p_i$ and $q = \sum_i y_i^2 q_i$, then $p + q \le 1$.

$$(\mathbf{x}^T \boldsymbol{A} \mathbf{x})(\mathbf{x}^T \boldsymbol{A}^{-1} \mathbf{x}) = (\lambda_1 p + \lambda_n q)(\lambda_1^{-1} p + \lambda_n^{-1} q) \le (p+q)^2 \frac{(\lambda_1 + \lambda_n)^2}{4\lambda_1\lambda_n} \le \frac{(\lambda_1 + \lambda_n)^2}{4\lambda_1\lambda_n}$$

after a few more steps of algebra involving geometry mean inequalities. ∎

6 TYING UP LOOSE ENDS

We are basically done. We use the Kantorovich inequality to bound the decrease. A bit more algebra yields the identity:

$$\frac{f(\mathbf{x}_{k+1})}{f(\mathbf{x}_k)} \le \left(\frac{\lambda_n - \lambda_1}{\lambda_n + \lambda_1}\right)^2 = \left(\frac{\kappa(\boldsymbol{Q}) - 1}{\kappa(\boldsymbol{Q}) + 1}\right)^2.$$