The quadratic penalty method for equality-constrained optimization

Mark S. Gockenbach

1 Introduction

The most straightforward methods for solving a constrained optimization problem convert it to a sequence of unconstrained problems whose solutions converge to the desired solution. The idea of the quadratic penalty method is to add to the objective function a term that penalizes infeasibility. The quadratic penalty function for

$$\begin{aligned} \min \quad & f(x) \\ s.t. \quad & g(x) = 0 \end{aligned} \tag{1}$$

is

$$Q(x;\mu) = f(x) + \frac{1}{2\mu}\|g(x)\|^2, \tag{2}$$

where μ is a positive scalar called the *penalty parameter*. Solving

$$\min_x Q(x;\mu)$$

yields a solution x_μ^*. Since smaller values of μ make any given infeasible point x less attractive when minimizing $Q(x;\mu)$, it is hoped that x_μ^* approaches the feasible set and hence that

$$x_\mu^* \to x^* \text{ as } \mu \to 0,$$

where x^* is a solution of (1).

A minimizer x_μ^* of $Q(\cdot;\mu)$ must satisfy

$$\nabla Q(x_\mu^*;\mu) = 0.$$

Since

$$\nabla Q(x_\mu^*;\mu) = \nabla f(x_\mu^*) + \frac{1}{\mu}\nabla g(x_\mu^*)g(x_\mu^*),$$

x_μ^* satisfies

$$\nabla f(x_\mu^*) = \nabla g(x_\mu^*)\left(-\frac{1}{\mu}g(x_\mu^*)\right).$$

If x^* is a minimizer of the NLP (1) and λ^* is the Lagrange multiplier corresponding to x^*, then

$$\nabla f(x^*) = \nabla g(x^*)\lambda^*.$$

These formulas suggest that, if $x_\mu^* \to x^*$, then

$$-\frac{1}{\mu}g(x_\mu^*) \to \lambda^* \text{ as } \mu \to 0.$$

This can be proved under certain mild conditions.

μ	x_μ^*	$\lvert g(x_\mu^*)\rvert$	$-\frac{1}{\mu}g(x_\mu^*)$
10^{-1}	$(0.34798, 1.0326)$	$1.8737 \cdot 10^{-1}$	-1.8737
10^{-2}	$(0.31586, 0.96015)$	$2.1660 \cdot 10^{-2}$	-2.1660
10^{-3}	$(0.31234, 0.95180)$	$3.4560 \cdot 10^{-3}$	-3.4560
10^{-4}	$(0.31164, 0.95038)$	$3.3540 \cdot 10^{-4}$	-3.3540
10^{-5}	$(0.31148, 0.95023)$	$2.2095 \cdot 10^{-5}$	-2.2095

Table 1: The results from Example 1.1.

Example 1.1 *As an example, I consider the NLP*

$$\begin{array}{ll} \min & f(x) \\ s.t. & g(x) = 0, \end{array}$$

where $f : \mathrm{R}^2 \to \mathrm{R}$ *and* $g : \mathrm{R}^2 \to \mathrm{R}$ *are defined by* $f(x) = (x_1 - 1)^2 + 2(x_2 - 2)^2$ *and* $g(x) = x_1^2 + x_2^2 - 1$, *respectively. The contours of* f *and the feasible set (the unit circle) are shown in Figure 1.*

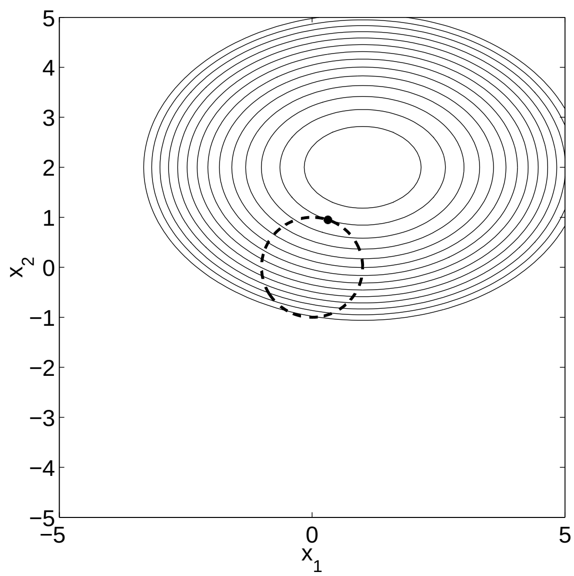

Figure 1: The contours of f and the feasible set determined by $g(x) = 0$ (see Example 1.1). The feasible set is the unit circle, shown as the dashed curve. The constrained minimizer is indicated by an asterisk.

A straightforward calculation shows that the global minimizer and corresponding Lagrange multiplier of the NLP are

$$x^* \doteq \begin{bmatrix} 0.31157 \\ 0.95022 \end{bmatrix}, \ \lambda^* \doteq -2.2095.$$

I applied the penalty method with $\mu = 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}$, *obtaining the results in Table 1. These results illustrate how the penalty method produces estimates of* x^* *and* λ^*.

One difficulty in achieving the results described above is that (1) and $Q(\cdot;\mu)$ can have many local minimizers. To develop a satisfactory theory, it is necessary to assume that the minimizers x_μ^* that are computed for various values of μ are consistent in that they are all related to the same solution x^* of (1). One way to do this is to consider a specific strict local minimizer x^* of (1) and to choose a neighborhood N of x^* in which x^* is the only minimizer of the NLP. I will prove below that, for μ sufficiently small, $Q(\cdot;\mu)$ has a unique minimizer in N, which will be designated as x_μ^*.

2 Analysis of the quadratic penalty method

Theorem 2.1 *Suppose that*

$$\begin{array}{ll} \min & f(x) \\ s.t. & g(x) = 0, \end{array}$$

where $f : \mathrm{R}^n \to \mathrm{R}$ and $g : \mathrm{R}^n \to \mathrm{R}^m$, has a strict local minimizer x^*, and let N be a neighborhood of x^* such that x^* is the unique solution of the NLP, subject to the further constraint that $x \in N$. Suppose further that x_μ^* is the unique minimizer of

$$Q(x;\mu) = f(x) + \frac{1}{2\mu}\|g(x)\|^2$$

over N for all μ sufficiently small, say $0 < \mu \le \hat{\mu}$. Then $0 < \mu_2 < \mu_1 \le \hat{\mu}$ implies

1. *$Q(x_{\mu_2}^*;\mu_2) \ge Q(x_{\mu_1}^*;\mu_1)$;*
2. *$\|g(x_{\mu_2}^*)\| \le \|g(x_{\mu_1}^*)\|$;*
3. *$f(x_{\mu_2}^*) \ge f(x_{\mu_1}^*)$.*

Proof: First of all,

$$\mu_2 < \mu_1 \;\Rightarrow\; Q(x_{\mu_2}^*;\mu_2) \ge Q(x_{\mu_2}^*;\mu_1) \ge Q(x_{\mu_1}^*;\mu_1),$$

which yields the first conclusion. Second, the optimality of $x_{\mu_1}^*$ and $x_{\mu_2}^*$ implies the inequalities

$$\begin{aligned} Q(x_{\mu_1}^*;\mu_1) &\le Q(x_{\mu_2}^*;\mu_1), \\ Q(x_{\mu_2}^*;\mu_2) &\le Q(x_{\mu_1}^*;\mu_2), \end{aligned}$$

that is,

$$\begin{aligned} f(x_{\mu_1}^*) + \frac{1}{2\mu_1}\|g(x_{\mu_1}^*)\|^2 &\le f(x_{\mu_2}^*) + \frac{1}{2\mu_1}\|g(x_{\mu_2}^*)\|^2, \\ f(x_{\mu_2}^*) + \frac{1}{2\mu_2}\|g(x_{\mu_2}^*)\|^2 &\le f(x_{\mu_1}^*) + \frac{1}{2\mu_2}\|g(x_{\mu_1}^*)\|^2. \end{aligned}$$

Adding these two inequalities and rearranging yields

$$\frac{1}{2}\left(\frac{1}{\mu_1} - \frac{1}{\mu_2}\right)\left(\|g(x_{\mu_1}^*)\|^2 - \|g(x_{\mu_2}^*)\|^2\right) \le 0.$$

Since $1/\mu_1 - 1/\mu_2 < 0$, this is possible only if

$$\|g(x_{\mu_1}^*)\| \ge \|g(x_{\mu_2}^*)\|,$$

which is the second conclusion. Finally,

$$\begin{aligned} & Q(x_{\mu_2}^*;\mu_1) \ge Q(x_{\mu_1}^*;\mu_1) \\ \Rightarrow\quad & f(x_{\mu_2}^*) + \frac{1}{2\mu_1}\|g(x_{\mu_2}^*)\|^2 \ge f(x_{\mu_1}^*) + \frac{1}{2\mu_1}\|g(x_{\mu_1}^*)\|^2 \\ \Rightarrow\quad & f(x_{\mu_2}^*) - f(x_{\mu_1}^*) \ge \frac{1}{2\mu_1}\left(\|g(x_{\mu_1}^*)\|^2 - \|g(x_{\mu_2}^*)\|^2\right) \ge 0, \end{aligned}$$

which yields the final conclusion. QED

3 The implicit function theorem

In order to prove that $Q(\cdot;\mu)$ has a solution for all $\mu > 0$ sufficiently small, I need to use the following theorem:

Theorem 3.1 (The implicit function theorem) *Suppose* $F : \mathrm{R}^n \times \mathrm{R}^k \to \mathrm{R}^n$ *is continuously differentiable and that* $x^* \in \mathrm{R}^n$, $y^* \in \mathrm{R}^k$ *satisfy*

$$F(x^*, y^*) = 0.$$

Let J^* *be the Jacobian of* F *with respect to* x*, evaluated at* (x^*, y^*)*. If* J^* *is nonsingular, then there exists a neighborhood* M *of* y^* *and a unique function* $f : M \to \mathrm{R}^n$ *such that* $f(y^*) = x^*$ *and*

$$F(f(y), y) = 0 \text{ for all } y \in M.$$

Moreover, f *is continuously differentiable.*

In other words, under the given conditions, it is possible to solve the system of n *equation for the first* n *variables in terms of the other* k *variables, at least locally around* y^**.*

4 Convergence of the quadratic penalty method

I wish to apply the implicit function theorem to the first-order optimality conditions of the NLP (1). This requires that I write the condition

$$\nabla Q(x;\mu) = 0 \tag{3}$$

in a form similar to the extended system (1). Introducing the variable λ, (3) is equivalent to

$$\begin{aligned} \nabla f(x) - \nabla g(x)\lambda &= 0, \\ -g(x) - \mu\lambda &= 0. \end{aligned}$$

I now define $F : \mathrm{R}^n \times \mathrm{R}^m \times \mathrm{R} \to \mathrm{R}^n \times \mathrm{R}^m$ by

$$F(x, \lambda, \mu) = \begin{bmatrix} \nabla f(x) - \nabla g(x)\lambda \\ -g(x) - \mu\lambda \end{bmatrix}.$$

If x^* is a local minimizer of NLP (1) and λ^* is the associated Lagrange multiplier, then

$$F(x^*, \lambda^*, 0) = 0.$$

By the implicit function theorem, I can solve for $F(x, \lambda, \mu) = 0$ for x^*, λ^* as functions of μ, provided the Jacobian of F with respect to (x, λ) is nonsingular at $(x^*, \lambda^*, 0)$.

The Jacobian of F with respect to (x, λ) is

$$\begin{bmatrix} \nabla^2 f(x) - \sum_{i=1}^m \lambda_i \nabla^2 g_i(x) & -\nabla g(x) \\ -\nabla g(x)^T & -\mu I \end{bmatrix} = \begin{bmatrix} \nabla^2 \ell(x;\lambda) & -\nabla g(x) \\ -\nabla g(x)^T & -\mu I \end{bmatrix}$$

(which is a symmetric matrix). At $(x^*, \lambda^*, 0)$, the Jacobian is

$$J^* = \begin{bmatrix} \nabla^2 \ell(x^*;\lambda^*) & -\nabla g(x^*) \\ -\nabla g(x^*)^T & 0 \end{bmatrix} \tag{4}$$

I will now show that J^* is nonsingular provided x^* is a regular point and x^*, λ^* satisfy the second-order sufficient condition

$$z \in \mathcal{N}(\nabla g(x^*)^T), z \neq 0 \;\Rightarrow\; z \cdot \nabla^2 \ell(x^*;\lambda^*) z > 0. \tag{5}$$

If $(z, w) \in \mathrm{R}^n \times \mathrm{R}^m$ satisfies $J^*(z, w) = 0$, then

$$\nabla^2 \ell(x^*; \lambda^*) z - \nabla g(x^*) w = 0, \tag{6}$$
$$\nabla g(x^*)^T z = 0. \tag{7}$$

Taking the dot product of both sides of (6) with z yields

$$\begin{aligned} z \cdot \nabla^2 \ell(x^*; \lambda^*) z - z \cdot \nabla g(x^*) w = 0 \quad &\Rightarrow \quad z \cdot \nabla^2 \ell(x^*; \lambda^*) z - (\nabla g(x^*)^T z) \cdot w = 0 \\ &\Rightarrow \quad z \cdot \nabla^2 \ell(x^*; \lambda^*) z = 0 \end{aligned}$$

(using the fact that $z \in \mathcal{N}(\nabla g(x^*)^T)$). But then

$$z \in \mathcal{N}(\nabla g(x^*)^T), \;\; z \cdot \nabla^2 \ell(x^*; \lambda^*) z = 0,$$

and therefore the second-order sufficient condition implies that $z = 0$. Equation (6) then implies that

$$\nabla g(x^*) w = 0,$$

which shows that $w = 0$ (since x^* is a regular point and hence the columns of $\nabla g(x^*)$ are linearly independent). Since the only vector (z, w) satisfying $J^*(z, w) = 0$ is the zero vector, J^* must be nonsingular. I will call x^* a *nonsingular point* if x^* is regular point that, together with λ^*, satisfies the second-order sufficient condition.

Assuming that x^* is a nonsingular point, the implicit function theorem yields continuous functions $x^* = x^*(\mu)$ and $\lambda^* = \lambda^*(\mu)$, defined for all μ sufficiently close to 0, such that

$$F(x^*(\mu), \lambda^*(\mu), \mu) = 0.$$

The continuity implies that $x^*(\mu) \to x^*$ and $\lambda^*(\mu) \to \lambda^*$ as $\mu \to 0$. By definition,

$$\lambda^*(\mu) = -\frac{1}{\mu} g(x^*(\mu)),$$

so

$$-\frac{1}{\mu} g(x^*(\mu)) \to \lambda^* \text{ as } \mu \to 0.$$

I have now proved the following theorem:

Theorem 4.1 *Suppose $f : \mathrm{R}^n \to \mathrm{R}$ and $g : \mathrm{R}^n \to \mathrm{R}^m$ are twice continuously differentiable and that x^* is local minimizer of*

$$\begin{aligned} \min \quad & f(x) \\ s.t. \quad & g(x) = 0 \end{aligned}$$

with Lagrange multiplier λ^*. If x^* is a nonsingular point, then there exists a neighborhood N of x^* such that, for all μ sufficiently small,

$$Q(x; \mu) = f(x) + \frac{1}{2\mu} \|g(x)\|^2$$

has a unique stationary point x^*_μ in N. Moreover, x^*_μ depends continuously on μ, with

$$x^*_\mu \to x^* \text{ as } \mu \to 0,$$

and

$$-\frac{1}{\mu} g(x^*_\mu) \to \lambda^* \text{ as } \mu \to 0.$$

It remains only to show that the stationary points x_μ^* in Theorem 4.1 are in fact local minimizers of $Q(\cdot;\mu)$. I begin by writing $\nabla^2 Q$ as follows:

$$\begin{aligned}\nabla^2 Q(x_\mu^*;\mu) &= \nabla^2 f(x_\mu^*) + \frac{1}{\mu}\nabla g(x_\mu^*)\nabla g(x_\mu^*)^T + \frac{1}{\mu}\sum_{i=1}^m g_i(x_\mu^*)\nabla^2 g_i(x_\mu^*) \\ &= \nabla^2 \ell(x_\mu^*;\lambda_\mu^*) + \frac{1}{\mu}\nabla g(x_\mu^*)\nabla g(x_\mu^*)^T,\end{aligned}$$

where

$$\lambda_\mu^* = -\frac{1}{\mu} g(x_\mu^*).$$

To prove that $\nabla^2 Q(x_\mu^*;\mu)$ is positive definite for all μ sufficiently small, I will argue by contradiction. If the desired conclusion does not hold, there exist sequences $\{\mu_k\}$ and $\{y^{(k)}\}$ such that

$$\begin{aligned}&\mu_k \to 0 \text{ as } k \to \infty, \\ &\|y^{(k)}\| = 1 \text{ for all } k, \\ &y^{(k)} \cdot \nabla^2 Q(x_{\mu_k}^*;\mu_k) y^{(k)} \le 0 \text{ for all } k.\end{aligned}$$

Since $\{y^{(k)}\}$ is bounded, a subsequence must converge to some $y \in \mathrm{R}^n$. Without loss of generality, I might as well assume that $y^{(k)} \to y$; then $\|y\| = 1$ because $\|y^{(k)}\| = 1$ for all k.

Since

$$\nabla^2 \ell(x_\mu^*;\lambda_\mu^*) \to \nabla^2 \ell(x^*;\lambda^*) \text{ as } \mu \to 0,$$

it follows that

$$y^{(k)} \cdot \nabla^2 \ell(x_{\mu_k}^*;\lambda_{\mu_k}^*) y^{(k)}$$

must be uniformly bounded, say

$$-M \le y^{(k)} \cdot \nabla^2 \ell(x_{\mu_k}^*;\lambda_{\mu_k}^*) y^{(k)} \le M \text{ for all } k.$$

It then follows that

$$0 \le \frac{1}{\mu_k} y^{(k)} \cdot \nabla g(x_{\mu_k}^*)\nabla g(x_{\mu_k}^*)^T y^{(k)} \le M \text{ for all } k,$$

that is,

$$0 \le y^{(k)} \cdot \nabla g(x_{\mu_k}^*)\nabla g(x_{\mu_k}^*)^T y^{(k)} \le \mu_k M \text{ for all } k.$$

This shows that

$$y^{(k)} \cdot \nabla g(x_{\mu_k}^*)\nabla g(x_{\mu_k}^*)^T y^{(k)} \to 0 \text{ as } k \to \infty.$$

But continuity implies that

$$y^{(k)} \cdot \nabla g(x_\mu^*)\nabla g(x_\mu^*)^T y^{(k)} \to y \cdot \nabla g(x^*)\nabla g(x^*)^T y \text{ as } k \to \infty.$$

Therefore

$$y \cdot \nabla g(x^*)\nabla g(x^*)^T y = 0$$

must hold, which implies that

$$\nabla g(x^*)^T y = 0$$

or, equivalently, that $y \in \mathcal{N}(\nabla g(x^*)^T)$.

It now follows, since x^* is by assumption a nonsingular point, that

$$y \cdot \nabla^2 \ell(x^*;\lambda^*) y > 0.$$

This implies that

$$y^{(k)} \cdot \nabla^2 \ell(x_{\mu_k}^*;\lambda_{\mu_k}^*) y^{(k)} > 0$$

for all k sufficiently large. But then

$$y^{(k)} \cdot \nabla^2 Q(x^*_{\mu_k}; \mu_k) y^{(k)} \geq y^{(k)} \cdot \nabla^2 \ell(x^*_{\mu_k}; \lambda^*_{\mu_k}) y^{(k)} > 0$$

for all k sufficiently large, a contradiction. This contradiction shows that $\nabla^2 Q(x^*_\mu; \mu)$ must be positive definite for all μ sufficiently small.

Theorem 4.1 can therefore be rephrased as follows.

Theorem 4.2 *Suppose $f : \mathrm{R}^n \to \mathrm{R}$ and $g : \mathrm{R}^n \to \mathrm{R}^m$ are twice continuously differentiable and that x^* is local minimizer of*

$$\begin{array}{ll} \min & f(x) \\ s.t. & g(x) = 0 \end{array}$$

with Lagrange multiplier λ^*. If x^* is a nonsingular point, then there exists a neighborhood N of x^* such that, for all μ sufficiently small,

$$Q(x; \mu) = f(x) + \frac{1}{2\mu} \|g(x)\|^2$$

has a unique minimizer x^*_μ in N. Moreover, x^*_μ depends continuously on μ, with

$$x^*_\mu \to x^* \text{ as } \mu \to 0,$$

and

$$-\frac{1}{\mu} g(x^*_\mu) \to \lambda^* \text{ as } \mu \to 0.$$

5 Ill-conditioning of the quadratic penalty method

In spite of the simplicity and strong convergence properties of the quadratic penalty method, it suffers from a significant disadvantage: The Hessian $Q(x^*_\mu; \mu)$ gets arbitrarily ill-conditioned as $\mu \to 0$, which implies that Newton's method and similar methods for minimizing $Q(\cdot; \mu)$ suffer from increasingly severe numerical difficulties as μ becomes small. Since it is necessary to take μ small in order to obtain convergence, this ill-conditioning is unavoidable.

The ill-conditioning of $\nabla^2 Q(x^*_\mu; \mu)$ is easy to demonstrate. It is necessary only to produce two unit vectors u, v such that

$$u \cdot Q(x^*_\mu; \mu) u \text{ and } v \cdot Q(x^*_\mu; \mu) v$$

differ greatly in magnitude. Since

$$\nabla^2 Q(x^*_\mu; \mu) = \nabla^2 \ell(x^*_\mu; \lambda^*_\mu) + \frac{1}{\mu} \nabla g(x^*_\mu) \nabla g(x^*_\mu)^T,$$

it suffices to take $u \in \mathcal{N}(\nabla g(x^*_\mu)^T)$ and $v \in \mathcal{R}(\nabla g(x^*_\mu)) = \mathcal{N}(\nabla g(x^*_\mu)^T)^\perp$.