Local convergence behaviour of the Newton method

Motivation

The Newton method is a second order method for the minimization of sufficiently smooth unconstrained objective functions, featuring in general a locally quadratic convergence speed. It is known for centuries and together with its variants is of an enormous importance for optimization [6]. For instance, Newton iterations are used as the elementary step in interior-point methods for constrained convex optimization [10]. Quasi-Newton methods such as BFGS use an estimation of the Hessian of the objective based on first order information from previous iterations and thus achieve an increased convergence speed with low memory requirements [1, 3, 4, 12]. Cubic regularization of the method greatly improves its global convergence properties [9].

For the pure Newton method, given by the iteration

$$x_{k+1} = x_k - (f''(x_k))^{-1} f'(x_k), \tag{1}$$

only local convergence results are available. If the starting point is sufficiently close to a non-degenerate minimum, $f''(x^*) \succ 0$, then the sequence of iterates will quadratically converge to the minimizer x^*, i.e., $||x_{k+1} - x^*|| \sim ||x_k - x^*||^2$. Outside of this quadratic convergence region little can be said about the behaviour of the method unless it is modified in some way.

Modifications which improve the global convergence behaviour can be made by adding regularization terms [7, 8, 9], decreasing the step-size, i.e., damping [11, 10], or confining the next iterate [2]. The performance of the different modifications will also depend on the function class to which they are applied.

Many of the modified methods feature a quadratic convergence speed if the iteration sequence gets near a non-degenerate local minimum of the objective function. However, an explicit check of whether the sequence of iterates has reached a quadratic convergence region of algorithm (1) has been designed only for the class of self-concordant functions [10], with the aim to apply it to interior-point methods for convex programming. We shall consider this class and the corresponding quadratic convergence condition in more detail in the next section.

Here we shall consider the above problem of checking whether a quadratic convergence region of algorithm (1) has been reached in application to other function classes. This will allow to augment a generic modified Newton method with a switch to the pure Newton iteration if the corresponding condition turns out to be validated. More precisely, we shall be concerned with the following question.

For a given function class $\mathcal{F}$, what are the conditions on the pair $(f'(x), f''(x))$ which ensure that x is in the quadratic convergence region of algorithm (1) for every function $f \in \mathcal{F}$?

This question has been investigated only for the class of self-concordant functions so far. We shall consider also two other function classes, namely the projectively self-concordant functions and functions with a Lipschitz-continuous Hessian. Below we explain these three cases in more detail.

Quadratic convergence region for self-concordant functions

The class of self-concordant functions is naturally linked to the Newton method by its affine invariance, and is hence well suited for minimization by this method. A standard reference on this topic is [10].

We consider the problem of minimizing a convex C^3 function $f : D \to \mathbb{R}$ defined on a convex domain D, with Hessian f'' positive definite everywhere on D. Given an iterate $x_k \in D$, the point x_{k+1} given by (1) can be interpreted as the minimizer of the strictly convex second order Taylor polynomial of f at x_k,

$$q_k(x) = f(x_k) + \langle f'(x_k), x - x_k\rangle + \frac{1}{2}(x - x_k)^T f''(x_k)(x - x_k). \tag{2}$$

The Newton algorithm is *affinely invariant*, i.e., its output does not change when computed in another coordinate system after an affine transformation. Therefore at x_k there exists no other natural norm than the Euclidean norm $||\cdot||_{x_k}$ defined by the Hessian $f''(x_k)$. In this norm the level subsets $\{x \,|\, q_k(x) \le c\}$ are norm balls around the minimizer x_{k+1} of the second order approximation $q_k(x)$. The current iterate lies at a distance

$$\rho_k = \sqrt{(x_{k+1} - x_k)^T f''(x_k)(x_{k+1} - x_k)} = \sqrt{f'(x_k)^T (f''(x_k))^{-1} f'(x_k)} \tag{3}$$

from the minimizer, which at the same time equals the norm of the gradient at the current iterate. This quantity, called *Newton decrement*, can also be expressed through the difference between the current function value and the minimum value of $q_k(x)$,

$$\begin{aligned} f(x_k) - q_k(x_{k+1}) &= -\langle f'(x_k), x_{k+1} - x_k\rangle - \frac{1}{2}(x_{k+1} - x_k)^T f''(x_k)(x_{k+1} - x_k) \\ &= \frac{1}{2} f'(x_k)^T (f''(x_k))^{-1} f'(x_k) = \frac{\rho_k^2}{2}. \end{aligned}$$

The situation at each iterate is hence characterized solely by the scalar ρ_k.

In order to be able to make assertions about the behaviour of the Newton algorithm, we must ensure that the norm defined by the Hessian f'' does not change too much when passing from the current iterate to the next. On the other hand, in an affinely invariant framework any such changes can be compared only against the step length measured in the norm defined by the current iterate. This leads to the following definition.

Definition 0.1. A convex C^3 function $f : D \to \mathbb{R}$ on a convex domain D is called *self-concordant* if it satisfies the inequality

$$|f'''(x)[h, h, h]| \le 2(f''(x)[h, h])^{3/2}$$

for all $x \in D$ and all tangent vectors h.

It is called *strongly* self-concordant if in addition $\lim_{x\to\partial D} f(x) = +\infty$.

Here the derivatives of f are as above treated as multi-linear maps on the space of tangent vectors, i.e.,

$$f''(x)[h, h] = \sum_{i,j=1}^{n} \frac{\partial^2 f(x)}{\partial x_i \partial x_j} h_i h_j, \qquad f'''(x)[h, h, h] = \sum_{i,j,k=1}^{n} \frac{\partial^3 f(x)}{\partial x_i \partial x_j \partial x_k} h_i h_j h_k,$$

where n is the dimension of the space. The exponent $\frac{3}{2}$ on the right-hand side has been introduced to obtain the same degree of homogeneity in h on both sides of the inequality.

If the current iterate is close to the minimum, i.e., if the Newton decrement ρ_k is small, then at the next step the distance to the minimum should approximately square. This means that ρ_{k+1} is upper bounded by a quantity of order $O(\rho_k^2)$. In order to understand whether the current iterate lies in the attraction basin of the minimum we need to quantity this upper bound on ρ_{k+1} as a function of ρ_k more precisely. We have the following result [10, Theorem 2.2.1].

Lemma 0.2. *For $\rho_k < 1$ the inequality $\rho_{k+1} \leq \frac{\rho_k^2}{(1-\rho_k)^2}$ holds.*

Corollary 0.3. *The Newton method converges from the initial point x_0 if the Newton decrement at this point satisfies $\rho_0 < \frac{3-\sqrt{5}}{2} \approx 0.3820$.*

Proof. It is easily seen that the relation in the corollary implies the relation $\rho_{k+1} < \rho_k$. □

The estimate in Lemma 0.2 is conservative. In one dimension we can easily calculate the following tight upper bound on ρ_{k+1}.

Lemma 0.4. *If the domain D has dimension 1, then for $0 < \rho_k < 1$ we have $\rho_{k+1} \leq 4 - \rho_k^2 - 4\sqrt{1-\rho_k^2}$ and this bound is achieved.*

Proof. We shall write the problem as an optimal control problem. By an affine change of the independent scalar variable x we may achieve $x_k = 0$, $f'(x_k) = -\rho_k$, $f''(x_k) = 1$, $x_{k+1} = \rho_k$.

Define the state variables $p = f'$, $h = f''$. Then $\rho_{k+1} = \frac{|p(\rho_k)|}{\sqrt{h(\rho_k)}}$, and the problem of finding an upper bound for ρ_{k+1} can be written as

$$\max \frac{|p(\rho_k)|}{\sqrt{h(\rho_k)}},$$

under the controlled dynamics

$$p' = h, \quad h' = 2uh^{3/2}, \qquad u \in [-1, 1],$$

and with initial conditions

$$p(0) = -\rho_k, \; h(0) = 1.$$

This is a classical optimal control problem which can be solved by the Pontryagin maximum principle. Introduce the variables ϕ, ψ, adjoint to p, h. Then the Pontryagin function is given by

$$\mathcal{H}(p, h, \phi, \psi, u) = \phi h + 2u\psi h^{3/2},$$

the optimal control is given by the maximizer $u^* = \operatorname{sgn}\psi$ of $\mathcal{H}$, and the adjoint variables satisfy the system

$$\phi' = -\frac{\partial \mathcal{H}}{\partial p} = 0, \quad \psi' = -\frac{\partial \mathcal{H}}{\partial h} = -\phi - 3u\psi\sqrt{h}.$$

It follows that $\phi = const$.

If at some point $\psi = 0$, then $\operatorname{sgn}\psi' = -\operatorname{sgn}\phi$, and hence ψ can change its sign at most once. We have two possibilities. Either first $u = 1$, and after the switch $u = -1$, or vice versa. In both cases the system can be integrated explicitly, with the switching time $x^* \in [0, \rho_k]$ being a parameter. We have

$$p(x) = \begin{cases} -\rho_k - \sigma + \frac{1}{\sigma - x}, & x \leq x^*, \\ -\rho_k - \sigma + \frac{2}{\sigma - x^*} - \frac{1}{\sigma - 2x^* + x}, & x \geq x^*, \end{cases} \qquad h(x) = \begin{cases} \frac{1}{(1-\sigma x)^2}, & x \leq x^*, \\ \frac{1}{(\sigma - 2x^* + x)^2}, & x \geq x^*. \end{cases}$$

Here $\sigma = \pm 1$ depending on which control is applied first.

Maximizing the objective function over σ and x^* yields the maximizer

$$(\sigma, x^*) = \left(1, 1 - \sqrt{\frac{1-\rho_k}{1+\rho_k}}\right)$$

with objective value $4 - \rho_k^2 - 4\sqrt{1-\rho_k^2}$. This proves our claim. □

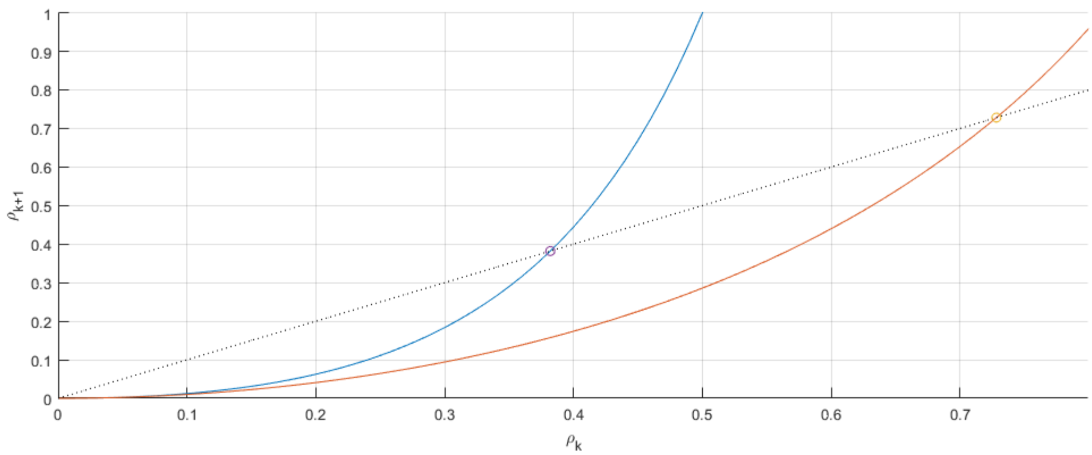

Figure 1: Upper bounds on ρ_{k+1}

Corollary 0.5. *If D is 1-dimensional, then the Newton algorithm converges with initial point x_0 if the Newton decrement at this point satisfies $\rho_0 < \zeta \approx 0.7282$, where ζ is the real root of the polynomial $\zeta^3 + 2\zeta^2 + 9\zeta - 8$.*

Proof. The value of ζ is easily obtained from the relation $4 - \zeta^2 - 4\sqrt{1-\zeta^2} = \zeta$. □

The upper bounds from Lemmas 0.2 and 0.4 are illustrated in Fig. 1. The exact upper bound for general dimension must lie in between these two curves.

Problem 0.1. *Find a tighter upper bound on ρ_{k+1} in dependence on ρ_k.*

The result is significant for the design of interior-point methods for solving constrained convex optimization problems. Here the constrained problem is turned into an unconstrained one by adding a smooth self-concordant barrier function to the linear objective. The weight of the barrier in the composite objective is adjusted at each step such that the current iterate always stays in the quadratic convergence region around the minimizer of the composite function. Increasing the radius of the convergence region will allow to make larger steps by choosing a more aggressive weight adjusting strategy. This in turn will lead to faster convergence to the optimum.

Quadratic convergence region for projectively self-concordant functions

The self-concordant functions as defined in [10] are not the only affinely invariant function class. A class with an even larger symmetry group is that of projectively self-concordant functions [5].

Definition 0.6. Let $D \subset \mathbb{R}^n$ be a convex domain. A *projectively self-concordant barrier* on D with parameter $\gamma \geq 0$ is a C^3 function $f : D \to \mathbb{R}$ satisfying

- $f''(x) - f'(x) \otimes f'(x) \succ 0$ for all $x \in D$,
- $\lim_{x \to \partial D} f(x) = +\infty$,
- $|f'''(x)[u,u,u] - 6f''(x)[u,u]f'(x)[u] + 4(f'(x)[u])^3| \leq 2\gamma(f''(x)[u,u] - (f'(x)[u])^2)^{3/2}$ for all $x \in D$, $u \in T_xD$.

Since the projectively self-concordant function will be used in a composite objective together with a linear function, the minimum x^* will be characterized by the relation $f'(x^*) = c$ for some fixed vector c. We shall solve this system of equations by a Newton-type algorithm.

Recall that the iteration step (1) can be reformulated by minimizing the quadratic approximation (2) around the current iterate x_k and taking the minimizer as the next iterate. In our setting, however, it is unnatural to consider the second order Taylor polynomial around x_k, because this function is not projectively self-concordant. We shall rather replace the Taylor polynomial by the function q_k satisfying the condition

$$q_k'''(x)[u,u,u] - 6q_k''(x)[u,u]q_k'(x)[u] + 4(q_k'(x)[u])^3 = 0 \tag{4}$$

for all $x \in D_k$ and $u \in T_x D_k$ with initial conditions $q_k'(x_k) = f'(x_k)$, $q_k''(x_k) = f''(x_k)$. Here D_k is the domain of definition of q_k. The next iterate x_{k+1} will then be the solution of the system of equations $q_k'(x) = c$. Both q_k and x_{k+1} can be computed explicitly, because the general solution of (4) is of the form $q(x) = -\frac{1}{2}\log Q(x)$, with Q a quadratic function.

Problem 0.2. *Express x_{k+1} explicitly as a function of x_k, $f'(x_k)$, $f''(x_k)$. Find conditions on the pair $(f'(x_0), f''(x_0))$ such that the iterative scheme outlined above with initial point x_0 converges to the solution of the system $f'(x) = c$.*

From the results in [5] it is to be expected that the convergence region will be larger than in the (affinely) self-concordant case from the previous section. Algorithms based on this class of functions are hence expected to be faster.

Quadratic convergence region for functions with Lipschitz-continuous Hessian

In this section we shall consider the Newton method on the class of functions with Lipschitz-continuous Hessian. For minimization of functions in this class Y.E. Nesterov and B.T. Polyak proposed to use a Newton method with cubic regularization. In [9, Theorem 3] it is shown that this modification is quadratically convergent near a non-degenerate local minimum. However, it is reasonable to expect that switching to the pure Newton method at the final stages of the iterative process is computationally less expensive and the convergence will be still faster. This motivates the following problem:

Problem 0.3. *Let $L > 0$ be a constant, and let $\mathcal{F}_{2,L}$ be the class of C^2 functions with Lipschitz-continuous Hessian, $||f''(x) - f''(y)|| \leq L||x - y||$. Find conditions on the pair $(f'(x), f''(x))$ ensuring that under the pure Newton iteration (1) the point x lies in the attraction basin of a local non-degenerate minimum x^* of f for every function f in the class $\mathcal{F}_{2,L}$.*

By possibly multiplying f with a positive constant, which does not change the sequence of Newton iterates, we may assume without loss of generality that $L = 1$.

References

[1] C.G. Broyden. The convergence of a class of double-rank minimization algorithms. *J. I. Math. Appl.*, 6:76–90, 1970.

[2] A.B. Conn, N.I.M. Gould, and P.L. Toint. *Trust region methods.* SIAM, Philadelphia, 2000.

[3] R. Fletcher. A new approach to variable metric algorithms. *Comput. J.*, 13:317–322, 1970.

[4] D. Goldfarb. A family of variable metric updates derived by variational means. *Math. Comput.*, 24:23–26, 1970.

[5] Roland Hildebrand. Projectively self-concordant barriers on convex sets. arXiv preprint 1909.01883, 2019.

[6] L.V. Kantorovich. Functional analysis and applied mathematics. *Uspechi Mat. Nauk*, 3(1):89–185, 1948.

[7] K. Levenberg. A method for the solutions of certain problems in least squares. *Quart. Appl. Math.*, 2:164–168, 1944.

[8] D. Marquardt. An algorithm for least-squares estimation of non-linear parameters. *SIAM J. Appl. Math.*, 11:431–441, 1963.

[9] Y.E. Nesterov and B.T. Polyak. Cubic regularization of Newton method and its global performance. *Math. Program. Ser. B*, 108(1):177–205, 2006.

[10] Yurii Nesterov and Arkadii Nemirovskii. *Interior-point Polynomial Algorithms in Convex Programming*, volume 13 of *SIAM Stud. Appl. Math.* SIAM, Philadelphia, 1994.

[11] J.M. Ortega and W.C. Rheinboldt. *Iterative solution of nonlinear equations in several variables.* Academic Press, New York, 1970.

[12] D.F. Shanno. Conditioning of quasi-newton methods for function minimization. *Math. Comput.*, 24:647–656, 1970.