Broyden Updating, the Good and the Bad!

Andreas Griewank

Abstract.

2010 Mathematics Subject Classification: 65H10, 49M99, 65F30
Keywords and Phrases: Quasi-Newton, secant condition, least change, bounded deterioration, superlinear convergence

> So far so good! We had an updating procedure (the 'full' secant method) that seemed to work provided that certain conditions of linear independence were satisfied, but the problem was that it did not work very well. In fact it proved to be quite numerically unstable.
> Charles Broyden in *On the discovery of the 'good Broyden' method* [6].

The idea of secant updating

As Joanna Maria Papakonstantinou recounted in her comprehensive historical survey [29], regula falsi and other variants of the secant method for solving one equation in one variable go back to the Babylonian and Egyptian civilizations nearly 4000 years ago. They may be viewed just as a poor man's version of what is now known as Newton's method, though we should also credit Al Tusi [20]. During antiquity the very concept of derivatives was in all likelihood unknown, and in modern times the evaluation (and in the multivariate case also factorization) of Jacobian matrices is frequently considered too tedious and computationally expensive.

The latter difficulty was certainly the concern of Charles Broyden in the sixties, when he tried to solve nonlinear systems that arose from the discretization of nonlinear reactor models for the English Electric Company in Leicester [6]. Now we know that, due to diffusion, the resulting system of ODEs must have been rather stiff, but that property was only identified and analyzed a few years later by Dahlquist. Nevertheless, Broyden and his colleagues already used some implicit time integration schemes, which required solving sequences of slightly perturbed nonlinear algebraic systems $F(x) = 0$ for $F : \mathbb{R}^n \mapsto \mathbb{R}^n$.

Broyden noted that one could avoid the effort of repeatedly evaluating and factoring the system Jacobian by exploiting secant information, i.e., function value differences

$$y_i \;\equiv\; F_i - F_{i-1} \quad \text{with} \quad F_j \equiv F(x_j) \quad \text{for} \quad j \le i$$

Here, $x_i \in \mathbb{R}^n$ denotes the current iterate and x_j, for $j < i$, disctinct points at which F has been evaluated previously. With $s_i \equiv x_i - x_{i-1}$ the new approximation B_i to the Jacobian $F'(x_i) \in \mathbb{R}^{n \times n}$

$$B_i s_i \;=\; y_i \;=\; F'(x_i)s_i + o(\|s_i\|) \tag{1}$$

The first order Taylor expansion on the right is valid if F has a Jacobian $F'(x) \in \mathbb{R}^{n\times n}$ that varies continuously in x. We will tacitly make this assumption throughout so that $F \in \mathcal{C}^1(\mathcal{D})$ on some open convex domain $\mathcal{D} \subset \mathbb{R}^n$ containing all evaluation points of interest.

In the univariate case of $n = 1$, one can divide by s_i to obtain $B_i = y_i/s_i \approx F'(x_i)$ uniquely. In the multivariate case, the secant condition merely imposes n conditions on the n^2 degrees of freedom in the new approximating Jacobian B_i. A natural idea is to remove the indeterminacy by simultaneously imposing earlier secant conditions $B_i s_j \;=\; y_j$, for $j = i - n + 1 \ldots i$. The resulting matrix equation for B_i has a unique solution provided the $n+1$ points x_{i-n+j}, for $j = 0 \ldots n$, are in *general position*, i.e., do not belong to a proper affine subspace of $\mathbb{R}^n$. Theoretically, that happens with probability 1, but in practice the step vectors s_j, for $j = i - n + 1 \ldots i$, are quite likely to be nearly linearly dependent, which leads to the observation of instability by Broyden cited above.

Rather than recomputing B_i from scratch, Broyden reasoned that the previous approximation B_{i-1} should be updated such that the current secant condition is satisfied, but $B_i v = B_{i-1} v$ in all directions $v \in \mathbb{R}^n$ orthogonal to s_i. As he found out 'after a little bit of scratching around', these conditions have the unique solution [2]

$$B_i \;=\; B_{i-1} + r_i s_i^\top \big/ s_i^\top s_i, \quad \text{with} \quad r_i \;\equiv\; y_i - B_{i-1}s_i \tag{2}$$

Here the outer product $C_i \equiv r_i s_i^\top / s_i^\top s_i$ of the column vector r_i and the row vector $s_i^\top$ represent a rank one matrix. This formula became known as the *good Broyden update*, because it seemed to yield better numerical performance than the so-called bad formula (6) discussed below. For a recent review of quasi-Newton methods see the survey by J. M. Martinez [25].

Broyden stated that the fact that $C_i = B_i - B_{i-1}$ turned out to be of rank one was *pure serendipity*. Even though he claimed *'When I was at University they did not teach matrices to physicists'*, he realized right away that the low rank property could be used to reduce the linear algebra effort for computing the next quasi-Newton step

$$s_{i+1} = -B_i^{-1} F_i$$

to $O(n^2)$. That compares very favourably with the $n^3/3$ arithmetic operations needed for a dense LU factorization of the new Jacobian $F'(x_i)$ to compute the Newton step $-F'(x_i)^{-1}F_i$. If the previous step is given by $s_i = -B_{i-1}^{-1}F_{i-1}$, one can easily check that the secant error vector r_i defined in (2) is identical to the new residual, i.e., $r_i = F_i$, which we will use below.

Tacking on a sequence of rank one corrections to an initial guess B_0, and reducing the linear algebra effort in the process looks more like an engineering trick than an algorithmic device of mathematical interest. Yet after a few years and in close collaboration with his coauthors John Dennis and Jorge Moré, a beautiful theory of superlinear convergence theory emerged [7], which was later built upon by other researchers and extended to many update formulas. For a much larger class of methods named after Charles Broyden and his coauthors Abbaffy and Spedicato, see [1].

Least change interpretation

John Dennis credits Jorge Moré with a short argument showing that the good Broyden formula is *a least change update*. Specifically, if we endow the real space of $n \times n$ matrices A with the inner product

$$\langle A, B\rangle \;\equiv\; Tr(A^\top B) = Tr(B^\top A)$$

then the corresponding norm

$$\|A\|_F \equiv \sqrt{\langle A, A\rangle} \;\geq\; \|A\| \tag{3}$$

is exactly the one introduced by Frobenius. It is bounded below by the consistent matrix norm $\|A\|$ induced by the Euclidean vector norm $\|v\|$ on $\mathbb{R}^n$. The affine variety

$$[y_i/s_i] \;\equiv\; \left\{B \in \mathbb{R}^{n\times n} : Bs_i = y_i\right\}$$

has the $n(n-1)$ dimensional tangent space $[0/s_i]$ and the n dimensional orthogonal complement

$$[0/s_i]^\perp \;\equiv\; \left\{vs_i^\top \in \mathbb{R}^{n\times n} : v \in \mathbb{R}^n\right\}$$

Hence, the smallest correction of B_{i-1} to obtain an element of $[y_i/s_i]$ is given by the correction

$$C_i \;=\; r_i s_i^\top / s_i^\top s_i \;\in\; [r_i/s_i] \cap [0/s_i]^\perp$$

For formal consistency we will set $C_i = 0$ if $s_i = 0 = y_i$, which may happen for all $i \geq j$ if we have finite termination, i.e., reach an iterate x_j with $F_j = 0$.

The geometry is displayed below and yields for any other element $A_i \in [y_i/s_i]$ by Pythagoras

$$\|B_{i-1} - A_i\|_F^2 - \|B_i - A_i\|_F^2 \;=\; \|C_i\|_F^2$$

In particular, we have the *nondeterioration property*

$$\|B_i - A_i\|_F \;\leq\; \|B_{i-1} - A_i\|_F$$

This to hold for all $A_i \in [y_i/s_i]$ is in fact equivalent to the least change property of the update. Broyden stated this property apparently for the first time in his survey paper [4], which he rarely cited afterwards. Moreover, nondeterioration can be equivalently stated in the operator norm as

$$\|B_i - A_i\| \;\leq\; \|B_{i-1} - A_i\| \tag{4}$$

which makes sense even on an infinite dimensional Hilbert space where $\|\cdot\|_F$ is undefined.

Sequential properties in the affine case

So far we have described the single least change update $C_i = r_i s_i^\top / s_i^\top s_i$, but the key question is of course how a sequence of them compound with each other. One can easily check that $B_{i+1} = B_i + C_{i+1} = B_{i-1} + C_i + C_{i+1}$ satisfies the previous secant condition $B_{i+1} s_i = y_i$ only if s_i and s_{i+1} are orthogonal so that $C_{i+1} s_i = 0$. In fact, exactly satisfying all n previous secant conditions is not even desirable, because that would lead back to the classical multivariate secant method, which was found to be rather unstable by Broyden and others. However, successive updates do not completely undo each other and thus eventually lead to good predictions $B_{i-1} s_i \approx y_i$.

Now we will briskly walk through the principal arguments for the case when F is affine on a finite dimensional Euclidean space. Later we will discuss whether and how the resulting relations extend to nonlinear systems and infinite dimensional Hilbert spaces. Suppose for a moment that our equation is in fact affine so that

$$F(x) = A\,x + b \quad \text{with} \quad A \in \mathbb{R}^{n\times n} \quad \text{and} \quad b \in \mathbb{R}^n.$$

Then the secant conditions over all possible steps $s_i = -B_{i-1}^{-1}F_{i-1}$ are satisfied by the exact Jacobian $A \in [y_i/s_i]$ since $y_i = A_i s_i$ by definition of F. Moreover, let us assume that A and all matrices B with $\|B-A\| \leq \|B_0-A\|$ have inverses with a uniform bound $\|B^{-1}\| \leq \gamma$. This holds by the Banach Perturbation Lemma [27] for all B_0 that are sufficiently close to a nonsingular A.

Then we can conclude, as Broyden did in [3], that all B_i are nonsingular and, consequently, all steps $s_i = -B_{i-1}^{-1}F_{i-1}$ are well defined and bounded by $\|s_i\| \leq \gamma\|F_{i-1}\|$. Repeatedly applying Pythagoras' identity we obtain for any i the telescoping result that

$$\sum_{j=1}^{i} \|C_j\|_F^2 \;=\; \|B_0 - A\|_F^2 - \|B_i - A\|_F^2 \;\leq\; \|B_0 - A\|_F^2.$$

Hence, we derive from $C_j s_j = r_j$ and the fact that the Frobenius norm is stronger than the operator norm that

$$\lim_j \|C_j\|_F \to 0 \quad \text{and} \quad \lim_j \|r_j\|/\|s_j\| \;\leq\; \lim_j \|C_j\| = 0. \tag{5}$$

Whereas these limits remain valid in the nonlinear case considered below, they hold in a trivial way in the affine case considered so far. This follows from the amazing result of Burmeister and Gay [12] who proved that Broyden's good method reaches the roots of affine equations exactly in at most $2\,n$ steps. The proof appears a little like an algebraic fluke and there is nothing monotonic about the approach to the solution. Moreover, the restriction that the ball with radius $\|B_0 - A\|$ contains no singular matrix can be removed by some special updating steps or line-searches as, for example, suggested in [26], [17], and [23], also for the nonlinear case.

The glory: Q-superlinear convergence

The property $\|r_j\|/\|s_j\| \to 0$ was introduced in [8] and is now generally known as the Dennis and Moré characterization of Q-superlinear convergence. The reason is that it implies, with our bound on the stepsize, that $\|r_j\|/\|F_{j-1}\| \leq \gamma^{-1}\|r_j\|/\|s_j\| \to 0$ and thus

$$\frac{\|F_{i+1}\|}{\|F_i\|} \to 0 \quad \Longleftrightarrow \quad \frac{\|x_{i+1} - x_*\|}{\|x_i - x_*\|} \;\to\; 0$$

The equivalence holds due to the assumed nonsingularity of A so that, in any pair of norms, the residual size $\|F(x)\|$ is bounded by a multiple of the distance

Charles Broyden and his fellow quasi-Newton musketeers, J. Dennis and J. Moré

$\|x - x_*\|$ and vice versa. Correspondingly, the central concept of Q-superlinear convergence is completely invariant with respect to the choice of norms, a highly desirable property that is not shared by the weaker property of Q-linear convergence, where the ratio of successive residual norms $\|F(x_j)\|$ or solution distances $\|x_i - x_*\|$ is merely bounded away from 1.

Under certain initial assumptions Q-superlinear convergence is also achieved in the nonlinear case, and under a compactness condition even in infinite dimensional space. All this without any exact derivative information or condition that the sequence of steps be in some sense linearly independent.

Originally, it was widely believed that to ensure superlinear convergence one had to establish the *consistency condition* that the B_i converge to the true Jacobian $F'(x_*)$. In fact, these matrices need not converge at all, but, theoretically, may wander around $F'(x_*)$ in a spiral, with the correction norms $\|C_j\|$ square summable but not summable. This means that the predicted increments $B_{i-1}s_i/\|s_i\|$ in the normalized directions $s_i/\|s_i\|$ cannot keep being substantially different from the actual increments $y_i/\|s_i\|$ because the $s_i/\|s_i\|$ belong to the unit sphere, which is compact in finite dimensions.

The seemingly counterintuitive nature of the superlinear convergence proof caused some consternation in the refereeing process for the seminal paper by Broyden, Dennis and Moré [7]. It eventually appeared in the IMA Journal of Applied Mathematics under the editorship of Mike Powell. Broyden had analyzed the affine case, John Dennis contributed the concept of bounded deterioration on nonlinear problems and Jorge Moré contributed the least change characterization w.r.t. the Frobenius norm leading to the proof of superlinear convergence. All this is not just for good Broyden, but for a large variety of unsymmetric and symmetric updates like BFGS, where the Frobenius norms must be weighted, which somewhat localizes and complicates the analysis.

More specifically, suppose one starts at x_0 in the vicinity of a root $x_* \in$

$F^{-1}(0)$ near which the Jacobian is nonsingular and Lipschitz continuous. Then the nondeterioration condition (4) becomes a bounded deterioration condition with A_i replaced by $F'(x_*)$ and a multiplicative factor $1+O(\|x_i - x_*\|)$ as well as an additive term $O(\|x_i - x_*\|)$ on the right-hand side. From that one can derive Q-linear convergence provided B_0 is close enough to $F'(x_*)$, which, in turn, implies Q-superlinear convergence by the perturbed telescoping argument. More generally, we have the chain of implications

Bounded deterioration
$\Longrightarrow$ Linear Convergence
$\Longrightarrow$ Q-superlinear Convergence.

Actually, R-linear convergence is enough for the second implication. This modularization of the analysis is a very strong point of the Broyden-Dennis-Moré framework [7] and has allowed many other researchers to communicate and contribute in an economical fashion.

Bad Broyden by inverse least change

The BDM mechanism also applies to so-called inverse updates, especially Broyden's second unsymmetric formula. It can be derived by applying the least change criterion to the approximating inverse Jacobian

$$H_i \;=\; B_i^{-1} \quad \text{with} \quad H_i\, y_i = s_i$$

The equation on the right is called the inverse secant condition, which must be satisfied by H_i if $B_i = H_i^{-1}$ is to satisfy the direct secant condition (1). After exchanging s_i and y_i and applying the good Broyden formula to H_i one obtains the inverse update on the left, which corresponds to the direct update of B_i on the right

$$H_i \;=\; H_{i-1} + \frac{(s_i - H_{i-1}y_i)y_i^\top}{y_i^\top y_i} \quad \Longleftrightarrow \quad B_i \;=\; B_{i-1} + \frac{r_i y_i^\top}{y_i^\top s_i} \tag{6}$$

The correspondence between the two representations can be derived from the so-called Sherman–Morrison–Woodbury formula [13] for inverses of matrices subject to low rank perturbations.

Broyden suggested this formula as well, but apparently he and others had less favourable numerical experience, which lead to the moniker *Bad Broyden update*. It is not clear whether this judgement is justified, since the formula has at least two nice features. First, the inverse is always well defined, whereas the inverse of the good Broyden update can be seen to blow up if $y_i^\top B_{i-1} s_i = 0$. Second, the bad Broyden update is invariant with respect to linear variable transformations in that applying it to the system $\tilde F(\tilde x) \equiv F(T\tilde x) = 0$ with $\det(T) \neq 0$ leads to a sequence of iterates $\tilde x_i$ related to the original ones by $x_i = T\tilde x_i$, provided one initializes $\tilde x_0 = T^{-1}x_0$ and $\tilde B_0 = B_0 T$. The good

Broyden formula, on the other hand, is dependent on the scaling of the variables via the Euclidean norm, but is independent of the scaling of the residuals, which strongly influences the bad Broyden formula. However, even for quasi-Newton methods based on the good Broyden update, the squared residual norm often enters through the back door, namely as merit function during a line-search. The resulting stabilized nonlinear equation solver is strongly affected by linear transformations on domain or range. In this brief survey we have only considered full step iterations and their local convergence properties.

Whether or not one should implement quasi-Newton methods by storing and manipulating the inverses H_i is a matter for debate. Originally, Broyden and his colleagues had apparently no qualms about this, but later it was widely recommended, e.g., by the Stanford school [14], that one should maintain a triangular factorization of the B_i for reasons of numerical stability. Now it transpires that the required numerical linear algebra games, e.g., chasing sub-diagonal entries, are rather slow on modern computer architectures. In any case, the trend is to limited memory implementations for large scale applications, in view of which we will first try to study the influence of the variable number n on Broyden updating.

Estimating the R-order and efficiency index

One might fault the property of Q-superlinear convergence for being not sufficiently discriminating, because it can be established for all halfway sensible updating methods. In view of the limiting case of operator equations on Hilbert spaces to be considered later, one may wonder how the convergence rate of quasi-Newton methods depends on the dimension n. A finer measure of how fast a certain sequence $x_i \to x_*$ convergences is the so-called R-order

$$\rho \;\equiv\; \liminf_i |\log \|x_i - x_*\||^{1/i}$$

The limit inferior on the right reduces to a proper limit when the sequence $x_i \to x_*$ satisfies $\|x_i - x_*\| \sim \|x_{i-1} - x_*\|^\rho$. This is well known to hold with $\rho = 2$ for all iterations generated by Newton's method from an x_0 close to a regular root x_*. Generally, the R-order [27] of a method is the infimum over ρ for all locally convergent sequences $(x_i)_{i=1\ldots\infty}$.

The result of Burmeister and Gay implies $2\,n$ step quadratic convergence of Broyden's good method on smooth nonlinear equations. That corresponds to an R-order of $\sqrt[2n]{2} = 1 + 1/(2n) + O(1/n^2)$. We may actually hope for just a little more by the following argument adapted from a rather early paper of Janina Jankowska [21]. Whenever a Jacobian approximation B_i is based solely on the function values $F_{i-j} = F(x_{i-j})$, for $j = 0\ldots n$, its discrepancy to the Jacobian $F'(x_*)$ is likely to be of order $O(\|x_{j-n} - x_*\|)$. Here we have assumed that things are going well in that the distances $\|x_i - x_*\|$ decrease monotonically towards 0, so that the function value at the oldest iterate x_{i-n} contaminates B_i most. Then the usual analysis of Newton-like iterations [9] yields the proportionality relation

$$\|x_{i+1} - x_*\| \ \sim \ \|x_{i-n} - x_*\| \, \|x_i - x_*\|$$

The first term on the right represents the error in the approximating Jacobian B_i multiplied by the current residual F_i of order $\|x_i - x_*\|$. Substituting the ansatz $\|x_i - x_*\| \sim c^{\rho^i}$ for some $c \in (0,1)$ into the recurrence and then taking the log base c one obtains immediately the relations

$$\rho^{i+1} \ \sim \ \rho^{i-n} + \rho^i \quad \Longrightarrow \quad 0 \ = \ P_n(\rho) \ \equiv \ \rho^{n+1} - 1 - \rho^n$$

Hence, we can conclude that the best R-order we may expect from Broyden updating is the unique positive root ρ_n of the polynomial $P_n(\rho)$.

For $n = 1$, both Broyden updating methods reduce to the classical secant scheme, which is well known [27] to have the convergence order $\rho_1 = (1+\sqrt{5})/2$. The larger n, the smaller ρ_n, and it was shown in [19] that asymptotically

$$P_n^{-1}(0) \ni \rho_n \ \approx \ 1 + \ln(n)/n \ \approx \ \sqrt[n]{n}$$

Here $a_n \approx b_n$ means that the ratio a_n/b_n tends to 1 as n goes to infinity. The second approximation means that we may hope for n step convergence of order n rather than just $2\,n$ step convergence of order 2 as suggested by the result of Burmeister and Gay.

The first approximation implies that the efficiency index [28] in the sense of Ostrowski (namely the logarithm of the R-order divided by the evaluation cost and linear algebra effort per step) satisfies asymptotically

$$\frac{\ln(\rho_n)}{OPS(F) + O(n^2)} \ \approx \ \frac{\ln(n)/n}{OPS(F) + O(n^2)} \ \geq \ \frac{\ln(2)}{n\,OPS(F) + O(n^3)}$$

The lower bound on the right-hand side represents Newton's method with divided difference approximation of the Jacobian, and dense refactorization at each iteration. As we can see there is a chance for Broyden updating to yield an efficiency index that is $\ln(n)/\ln(2) = \log_2 n$ times larger than for Newton's method under similar conditions.

This hope may not be in vain since it was shown in [19] that the R-order ρ_n is definitely achieved when the Jacobian is updated by the *adjoint Broyden* formula

$$B_i = B_{i-1} + r_i r_i^\top (F'(x_i) - B_{i-1}) / \, r_i^\top r_i$$

However, this rank-one-update is at least twice as expensive to implement since it involves the transposed product $F'(x_i)^\top r_i$, which can be evaluated in the reverse mode of Algorithmic Differentiation. The latter may be three times as expensive as pure function evaluation, so that the efficiency gain on Newton's method can be bounded below by $(\log_2 n)/4 = \log_{16} n$.

Whether or not simple Broyden updating itself achieves the optimal R-order ρ_n has apparently not yet been investigated carefully. To be fair, it should be noted that taking roughly $n/\log(n)$ simplified Newton steps before reevaluating and refactorizing the Jacobian in the style of Shamanskiĭ [22], yields the convergence order near $1+n/\log(n)$ for any such cycle and the corresponding effort is approximately $[n\,OPS(F)+O(n^3)][1+1/\log(n)]$. The resulting efficiency index is asymptotically identical to the optimistic estimate for Broyden updating derived above.

Pushing n to infinity

While Broyden updating is well established in codes for small and medium scale problems, its usefulness for large dimensional problems is generally in doubt. The first author who applied and analyzed Broyden's method to a control problem in Hilbert space was Ragnar Winther [31]. Formally, it is easy to extend the Broyden method to an operator equation $y = F(x) = 0$ between a pair of Hilbert spaces X and Y. One simply has to interpret transposition as taking the adjoint so that $v^\top$ represents a linear function in $X = X^*$ such that $v^\top w \equiv \langle v, w\rangle$ yields the inner product. The Good Broyden Update is still uniquely characterized by the nondeterioration condition (4) in terms of the operator norm $\|\cdot\|$. This implies bounded nondeterioration in the nonlinear case and everything needed to derive local and linear convergence goes through.

However, the least change characterization and its consequences cannot be extended, because there is no generalization of the Frobenius norm (3) and the underlying inner product to the space $\mathcal{B}(X,Y)$ of bounded linear operators. To see this, we simply have to note that, in n dimensions, the Frobenius norm of the identity operator is n, the sum of its eigenvalues. That sum would be infinite for the identity on l^2, the space of square summable sequences to which all separable Hilbert spaces are isomorphic. There is apparently also no other inner product norm on $\mathcal{B}(X,Y)$ that is at least as strong as the operator norm so that the implication (5) would work.

These are not just technical problems in extending the superlinear result, since X is infinite dimensional exactly when the unit ball and, equivalently, its boundary, the unit sphere, are not compact. That means one can keep generating unit directions $\bar{s}_i \equiv s_i/\|s_i\|$ along which the current approximation B_i is quite wrong. Such an example with an orthogonal sequence of s_i was given by Griewank [18]. There, on an affine bicontinuous problem, Broyden's method with full steps converges only linearly or not at all.

To derive the basic properties of Broyden's method in Hilbert space we consider an affine equation $0 = F(x) \equiv A\,x - b$ with a bounded invertible operator $A \in \mathcal{B}(Y,X)$. Then we have the discrepancies

$$D_i \;=\; A^{-1}B_i - I \in \mathcal{B}(X,Y) \quad \text{and} \quad E_i \;\equiv\; D_i^\top D_i \in \mathcal{B}(X)$$

where $D_i^\top \in \mathcal{B}(Y,X)$ denotes the adjoint operator to D_i and we abbreviate $\mathcal{B}(X) \equiv \mathcal{B}(X,X)$ as usual. By definition, E_i is selfadjoint and positive semidefinite. Now the Broyden good update can be rewritten as

$$D_{i+1} \ = \ D_i \left(I - \bar{s}_i \bar{s}_i^\top\right) \quad \Longrightarrow \quad E_{i+1} \ \equiv \ E_i - \bar{r}_i \bar{r}_i$$

with $\bar{r}_i \equiv A^{-1} r_i / \|s_i\|$.

In the finite dimensional case one could show that the Frobenius norm of the D_i decreases monotonically. Now we see that the operators E_i are obtained from the $E_0 = D_0^\top D_0$ by the consistent subtraction of rank-one terms. Hence, they have a selfadjoint semidefinite limit E_*. This implies, by a generalization of the interlacing eigenvalue theorem, that the eigenvalues $(\lambda_j(E_i))_{j=1\ldots\infty}$ of E_i are monotonically declining towards their limits $(\lambda_j(E_*))_{j=1\ldots\infty}$. Correspondingly, we find for the singular values $\sigma_j(D_i) = \sqrt{\lambda_j(E_i)}$ of the D_i that

$$\sigma_j(D_{i+1}) \ \le \ \sigma_j(D_i) \quad \text{and} \quad \sigma_j(D_i) \to \sqrt{\lambda_j(E_*)} \quad \text{for} \quad i \to \infty$$

Similarly, it was proven by Fletcher that the BFGS update monotonically moves all eigenvalues of the symmetric discrepancy $B_*^{-1/2} B_i B_*^{-1/2} - I$ between the Hessian B_* and its approximations B_i towards zero. With regards to convergence speed it was shown in [18] for $C^{1,1}$ operator equations that Broyden's method yields locally

$$\limsup_{i \to \infty} \|A^{-1} F_{i+1}\| / \|A^{-1} F_i\| \ \le \ \sigma_\infty(D_0) \ \equiv \ \lim_{j \to \infty} \sigma_j(D_0)$$

In other words, the Q-factor is bounded by the essential spectrum $\sigma_\infty(D_0)$ of the initial relative discrepancy $D_0 = A^{-1} B_0 - I$. Hence, we must have Q-superlinear convergence if D_0 or, equivalently, just $B_0 - A$ is compact, an assumption that is of course trivial in finite dimensions. Equivalently, we can require the preconditioned discrepancy D_0 to be compact or at least to have a small essential norm. Thus we can conclude that Broyden updating will yield reasonable convergence speed in Hilbert space if D_0 is compact or has at least a small essential norm $\sigma_\infty(D_0) = \sigma_\infty(D_j)$. It is well known that the essential norm is unaffected by modifications of finite rank. On the other hand, all singular values $\sigma_j(D_0) > \sigma_\infty(D_0)$ are effectively taken out as far as the final rate of convergence is concerned.

Limited memory and data sparse

For symmetric problems the idea of limited memory approximations to the Hessian of the objective [24] has been a roaring success. In the unsymmetric case things are not so clear. Whereas in the unconstrained, quadratic optimization case conjugate gradients generates the same iterates as BFGS in an almost memoryless way, there is, according to a result of Faber and Manteuffel [11], no short recurrence for unsymmetric real problems. Correspondingly, the more or less generic iterative solver GMRES for linear problems requires $2\,i$ vectors of storage for its first i iterations. The same appeared be true of Broyden's method, where starting from a usually diagonal B_0, one could store the secant pairs (s_j, y_j) for $j = 1 \dots i$.

The same appeared be true for Broyden's method in inverse form, where starting from an usually diagonal $H_0 = B_0^{-1}$ one could store the secant pairs (s_j, z_j) with $z_j \equiv H_{j-1} y_j$ for $j = 1 \dots i$. Then the inverse Hessian approximations have the product representation

$$H_i = \left[H_{i-1} + \frac{(s_i - z_i) s_i^\top H_{i-1}}{s_i^\top z_i} \right] = \prod_{j=1}^{i} \left[I + \frac{(s_j - z_j) s_j^\top}{s_j^\top z_j} \right] \; H_0$$

Deuflhard et al. noticed in [10] that for the fullstep iteration successive s_j and $s_{j+1} = -H_j F_j$ satisfy the relation $s_{j+1} = (s_j - z_j) \|s_j\|^2 / s_i^\top z_j$. Hence, one only needs to store the s_j and one can then cheaply reconstruct the z_j for applying the inverse in product form to any vector v usually the current residual F_i. Hence the storage requirement is only $i + O(1)$ vectors of length n up to the i-th iteration. In contrast the storage requirement for i iterations of Bad Broyden appears to be twice as large [10], so at least in that sense the derogatory naming convention is justified. In either case, one normally wishes to limit the number of vectors to be stored a priori and thus one has to develop strategies for identifying and discarding old information. This issue has been extensively studied for the limited memory BFGS method and for Broyden updating it has been the focus of a recent PhD thesis [30]. Usually one wishes to get rid of information from earlier iterates because nonlinearity may render it irrelevant or even misleading near the current iterates. On discretizations of infinite dimensional problems, one may wish to discard all corrections of a size close to the essential norm $\sigma_\infty(D_0)$, since no amount of updating can reduce that threshhold.

In good Broyden updating the correction made to any row of the approximating Jacobian is completely independent of what goes on in the other rows. In other words we are really updating the gradients ∇F_k of the component functions F_k independently. That shows immediately that one can easily use the method for approximating rectangular Jacobians $F'(x)$ for $F : \mathbb{R}^n \mapsto \mathbb{R}^m$ with m independent of n. Also in updating the $k-$th row one can disregard all variables that have no impact on F_k so that the corresponding Jacobian entries are zero. The resulting sparse update is known as Schubert's method [5]. The least change characterization now applies in the linear subspace of matrices with the appropriate sparsity pattern, and the whole BDM locally linear and Q-superlinear convergence goes through without any modification. However, since the update matrices C_j are now of high rank, there is no longer any advantage compared to Newton's method with regards to the linear algebra effort per step.

On the other hand, large sparse Jacobians can often be evaluated exactly, possibly using algorithmic differentiation [16], at an entirely reasonable cost. In particular it was found that none of the constraint Jacobians in the optimization test collection CUTEr takes more than 18 times the effort of evaluating the vector functions of constraints themselves. Since the sparsity patterns also tend to be quite regular, no methods based on Broyden type updating [15] can here compete with methods based on exact derivatives values.

Whether or not that situation is really representative for problems from applications is not entirely clear.

In any case we have to count the inability to effectively exploit sparsity as part of the Bad about Broyden updating. Still, there is a lot of Good as well, for which we have to thank primarily Charles Broyden, who passed away last year at the age of 78 after an eventful life with various professional roles and countries of residence.

Acknowledgement. The author is indebted to Jorge Moré, Trond Steihaug, and other colleagues for discussions on the historical record.

References

[1] Joszef Abaffy, Charles Broyden, and Emilio Spedicato. A class of direct methods for linear systems. *Numer. Math.*, 45(3):361–376, 1984.

[2] C. G. Broyden. A class of methods for solving nonlinear simultaneous equations. *Math. Comp.*, 19:577–593, 1965.

[3] C. G. Broyden. The convergence of single-rank quasi-Newton methods. *Math. Comp.*, 24:365–382, 1970.

[4] C. G. Broyden. Recent developments in solving nonlinear algebraic systems. In *Numerical methods for nonlinear algebraic equations (Proc. Conf., Univ. Essex, Colchester, 1969)*, pages 61–73. Gordon and Breach, London, 1970.

[5] C. G. Broyden. The convergence of an algorithm for solving sparse nonlinear systems. *Math. Comp.*, 25:285–294, 1971.

[6] C. G. Broyden. On the discovery of the "good Broyden" method. *Math. Program.*, 87(2, Ser. B):209–213, 2000. Studies in algorithmic optimization.

[7] C. G. Broyden, J. E. Jr. Dennis, and J. J. Moré. On the local and superlinear convergence of quasi-Newton methods. *JIMA*, 12:223–246, 1973.

[8] J. E. Dennis, Jr. and Jorge J. Moré. A characterization of superlinear convergence and its application to quasi-Newton methods. *Math. Comp.*, 28:549–560, 1974.

[9] J. E. Jr. Dennis and R. B. Schnabel. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations.* Prentice-Hall, 1996.

[10] Peter Deuflhard, Roland Freund, and Artur Walter. Fast secant methods for the iterative solution of large nonsymmetric linear systems. In *IMPACT of Computing in Science and Engineering*, pages 244–276, 1990.

[11] Vance Faber and Thomas Manteuffel. Necessary and sufficient conditions for the existence of a conjugate gradient method. *SIAM J. Numer. Anal.*, 21(2):352–362, 1984.

[12] D. M. Gay. Some convergence properties of Broyden's method. *SIAM J. Numer. Anal.*, 16:623–630, 1979.

[13] D. M. Gay and R. B. Schnabel. Solving systems of nonlinear equations by Broyden's method with projected updates. *In Nonlinear Programming 3, O. Mangasarian, R. Meyer and S. Robinson, eds., Academic Press, NY*, pages 245–281, 1978.

[14] Philip E. Gill, Walter Murray, and Margaret H. Wright. *Numerical linear algebra and optimization. Vol. 1.* Addison-Wesley Publishing Company Advanced Book Program, Redwood City, CA, 1991.

[15] A. Griewank and A. Walther. On constrained optimization by adjoint based quasi-Newton methods. *Opt. Meth. and Soft.*, 17:869–889, 2002.

[16] A. Griewank and A. Walther. *Principles and Techniques of Algorithmic Differentiation, Second Edition.* SIAM, 2008.

[17] Andreas Griewank. The "global" convergence of Broyden-like methods with a suitable line search. *J. Austral. Math. Soc. Ser. B*, 28(1):75–92, 1986.

[18] Andreas Griewank. The local convergence of Broyden-like methods on Lipschitzian problems in Hilbert spaces. *SIAM J. Numer. Anal.*, 24(3):684–705, 1987.

[19] Andreas Griewank, Sebastian Schlenkrich, and Andrea Walther. Optimal r-order of an adjoint Broyden method without the assumption of linearly independent steps. *Optim. Methods Softw.*, 23(2):215–225, 2008.

[20] Hermann Hammer and Kerstin Dambach. Sharaf al-tusi, ein vorläufer von newton und leibnitz. *Der mathematische und naturwissenschaftliche Unterricht*, 55(8):485–489, 2002.

[21] Janina Jankowska. Theory of multivariate secant methods. *SIAM J. Numer. Anal.*, 16(4):547–562, 1979.

[22] C. T. Kelley. A Shamanskiĭ-like acceleration scheme for nonlinear equations at singular roots. *Math. Comp.*, 47(176):609–623, 1986.

[23] Dong-Hui Li and Masao Fukushima. A derivative-free line search and global convergence of Broyden-like method for nonlinear equations. *Optim. Methods Softw.*, 13(3):181–201, 2000.

[24] Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Math. Programming*, 45(3, (Ser. B)):503–528, 1989.

[25] José Mario Martínez. Practical quasi-Newton methods for solving nonlinear systems. *J. Comput. Appl. Math.*, 124(1–2):97–121, 2000. Numerical analysis 2000, Vol. IV, Optimization and nonlinear equations.

[26] J. J. Moré and J. A. Trangenstein. On the global convergence of Broyden's method. *Math. Comp.*, 30(135):523–540, 1976.

[27] J. M. Ortega and W. C. Reinboldt. Iterative Solution of Nonlinear Equations in Several Variables. *Academic Press*, 2000.

[28] A. Ostrowski. *Solution of Equations and Systems of Equations*. Academic Press, New York, 1966.

[29] Joanna Maria Papakonstantinou. *Historical Development of the BFGS Secant Method and Its Characterization Properties*. PhD thesis, Rice University, Houston, 2009.

[30] Bart van de Rotten. *A limited memory Broyden method to solve high-dimensional systems of nonlinear equations*. PhD thesis, Mathematisch Instituut, Universiteit Leiden, The Netherlands, 2003.

[31] Ragnar Winther. *A numerical Galerkin method for a parabolic control problem*. PhD thesis, Cornell University, 1977.

Andreas Griewank
Institut für Mathematik
Humboldt Universität zu Berlin
Unter den Linden 6
10099 Berlin
Germany
griewank@mathematik.hu-berlin.de