



HAL
open science

Automatic differentiation and the step computation in the limited memory BFGS method

Jean Charles Gilbert, Jorge Nocedal

► **To cite this version:**

Jean Charles Gilbert, Jorge Nocedal. Automatic differentiation and the step computation in the limited memory BFGS method. Applied Mathematics Letters, 1993, 6 (3), pp.47-50. 10.1016/0893-9659(93)90032-I . hal-04140351

HAL Id: hal-04140351

<https://inria.hal.science/hal-04140351>

Submitted on 25 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution - NoDerivatives 4.0 International License

AUTOMATIC DIFFERENTIATION AND THE STEP COMPUTATION IN THE LIMITED MEMORY BFGS METHOD

Jean Charles Gilbert¹ and Jorge Nocedal²

Abstract. It is shown that the two-loop recursion for computing the search direction of a limited memory method for optimization can be derived by means of the reverse mode of automatic differentiation applied to an auxiliary function.

1. Introduction

The problem of finding efficient procedures for automatically calculating the gradient of a function has recently received much attention [4]. It is known that the *reverse mode* of automatic differentiation can provide the value of the gradient at a cost that is not much greater than that required to evaluate the function.

On the other hand the limited memory BFGS method – an optimization method designed for very large unstructured problems – has recently become popular. This method attempts to mimic the very successful BFGS variable metric method, but without storing matrices. To do this it saves several pairs of vectors $\{y_i, s_i\}$, $i = 1, \dots, m$ that implicitly define the iteration matrix H . The search direction of the limited memory method is then computed by $d = -Hg$, where H can be viewed as an approximation of the inverse Hessian of the objective function at the current iterate and g is the gradient of this function. Since the matrix H is not formed, but is only represented implicitly by the set of vectors $\{y_i, s_i\}$, a formula [6] is required to calculate the product Hg directly from the vectors $\{y_i, s_i\}$ and the gradient g . It turns out that this formula is not unique; one can devise various equivalent expressions, some of which are more economical than others [1]. For unconstrained optimization the most efficient formula for computing d , given in [6], consists of a two-loop recursion involving the vectors $\{y_i, s_i\}$ and the gradient g . In this paper we show that this two-loop recursion can be viewed as an application of the reverse mode of automatic differentiation. Thus, we find a connection between two apparently unrelated subjects: the step computation in a limited memory method and automatic differentiation.

2. Adjoint code of a program computing a function

Suppose that a function

$$f : x = (x_1, \dots, x_n) \in \mathbb{R}^n \longrightarrow f(x) \in \mathbb{R}$$

¹ INRIA, Domaine de Voluceau, Rocquencourt, B.P. 105, 78153 Le Chesnay Cedex (France).

² Department of Electrical Engineering and Computer Science, Northwestern University, Evanston IL 60208. This author was supported by the National Science Foundation Grant INT-9101901, and by the Department of Energy Grant DE-FG02-87ER25047.

is computed by a program executing the following sequence of instructions

$$\begin{aligned} & \mathbf{for} \ k := 1 \ \mathbf{to} \ K \ \mathbf{do} \ x_{\mu_k} := \varphi_k(\{x_{P_k}\}); \\ & f := x_N. \end{aligned} \tag{2.1}$$

For convenience, we have denoted by x_1, \dots, x_n the *independent variables*, with respect to which the gradient of f is desired, and by x_{n+1}, \dots, x_N ($N \geq n$) all the other variables used in the program. There is, however, no meaningful ordering in this notation. Instruction k in (2.1) modifies the variable x_{μ_k} , $\mu_k \in \{1, \dots, N\}$, by using an intermediate function φ_k depending on $\{x_{P_k}\}$, which is the collection of variables x_j with indices j in some subset $P_k \subset \{1, \dots, N\}$. After assigning a value to the variable x , this program will provide the value of $f(x)$ in the variable x_N .

It has been shown [8] that the gradient $\nabla f(x)$ of f at a given point x can be evaluated by first executing the original program (2.1), storing some partial derivatives $\partial\varphi_k/\partial x_j$ for $j \in P_k$, and then executing the following *adjoint code*:

$$\begin{aligned} & \mathbf{for} \ i := 1 \ \mathbf{to} \ N - 1 \ \mathbf{do} \ \bar{x}_i := 0; \\ & \bar{x}_N := 1; \\ & \mathbf{for} \ k := K \ \mathbf{down\ to} \ 1 \ \mathbf{do} \ \{ \\ & \quad \bar{x}_i := \bar{x}_i + \frac{\partial\varphi_k}{\partial x_i} \bar{x}_{\mu_k}, \ \forall i \in P_k \setminus \{\mu_k\}; \\ & \quad \bar{x}_{\mu_k} := \frac{\partial\varphi_k}{\partial x_{\mu_k}} \bar{x}_{\mu_k}; \} \\ & \mathbf{for} \ i := 1 \ \mathbf{to} \ n \ \mathbf{do} \ \frac{\partial f}{\partial x_i} := \bar{x}_i. \end{aligned} \tag{2.2}$$

Here we used the notation $A \setminus B$ to denote the set of elements that belong to A but not to B . The variable \bar{x}_i is called the *adjoint variable* associated to x_i . Its value is the current evaluation of the derivative of f with respect to x_i . This technique is also called the *reverse mode* of automatic differentiation; see [7], [3], [5] or [2] for an introduction to the subject. Its main advantage is that it computes $f(x)$ and its gradient $\nabla f(x)$ in a time $T(f, \nabla f)$ satisfying

$$T(f, \nabla f) \leq C T(f), \tag{2.3}$$

where C is some constant and $T(f)$ is the time to compute $f(x)$ by algorithm (2.1). When the functions φ_k are restricted to be the ones available in FORTRAN, it is reasonable to bound C by 5, see [3].

3. Application to the L-BFGS method

In a typical iteration of the limited memory BFGS method (L-BFGS), one is given a symmetric and positive definite matrix H_0 and m pairs of vectors $\{y_0, s_0\}, \dots, \{y_{m-1}, s_{m-1}\}$, where m is some integer. Each of these pairs satisfies the condition $\rho_i \equiv (y_i^\top s_i)^{-1} > 0$. The matrix H_0 is then updated m times using the BFGS formula, i.e., for $i = 0, \dots, m-1$:

$$H_{i+1} = V_i^\top H_i V_i + \rho_i s_i s_i^\top, \tag{3.1}$$

where

$$V_i = I - \rho_i y_i s_i^\top; \tag{3.2}$$

see [6]. The resulting matrix, H_m , is then used to compute the search direction

$$d = -H_m g,$$

where g is the current value of the gradient of the objective function.

To interpret this step computation in terms of automatic differentiation, we need to express $H_m g$ as the gradient of a real-valued function. A natural candidate for this is

$$f_m(g) = \frac{1}{2} g^\top H_m g,$$

since, due to the symmetry of H_m , $\nabla_g f_m(g)$ is precisely $H_m g$. From (3.1) we find that $f_m(\cdot)$ can easily be expressed in terms of $f_{m-1}(\cdot)$:

$$f_m(g) = f_{m-1}(V_{m-1} g) + \frac{\rho_{m-1}}{2} (s_{m-1}^\top g)^2.$$

Therefore, if we define

$$q_m = g, \quad q_{i-1} = V_{i-1} q_i \quad \text{for } i = m, \dots, 1, \quad (3.3)$$

we find by induction that

$$f_m(g) = f_0(q_0) + \sum_{i=0}^{m-1} \frac{\rho_i}{2} (s_i^\top q_{i+1})^2.$$

Using this formula, the computation of $f_m(g)$ can be performed by the following algorithm. We assume that the scalars ρ_i have been computed beforehand, and to save storage, we place all the q_i in the same vector q . Recall that the vectors q_i are updated by (3.3), where the matrices V_i are defined by (3.2).

$$\begin{aligned} & f := 0; \\ & q := g; \\ & \mathbf{for } i := m - 1 \mathbf{ down to } 0 \mathbf{ do } \{ \\ & \quad \alpha_i := s_i^\top q; \\ & \quad f := f + \frac{\rho_i}{2} \alpha_i^2; \\ & \quad q := q - \rho_i \alpha_i y_i; \} \\ & f := f + \frac{1}{2} q^\top H_0 q. \end{aligned} \quad (3.4)$$

The product $H_m g$ will now be computed by the adjoint code of (3.4). Let \bar{f} , \bar{q} , $\bar{\alpha}_0$, \dots , $\bar{\alpha}_{m-1}$ be the adjoint variables corresponding to f , q , α_0 , \dots , α_{m-1} . The adjoint code is obtained by writing the adjoint instructions corresponding to the instructions in (3.4), in the reverse order of execution. After the initialization of the adjoint variables,

$$\bar{f} := 1; \quad \bar{q} := 0; \quad \bar{\alpha}_0 := 0; \dots; \bar{\alpha}_{m-1} := 0,$$

the code continues as follows:

$$\begin{aligned}
& \bar{q} := H_0 q; \\
& \mathbf{for} \ i := 0 \ \mathbf{to} \ m - 1 \ \mathbf{do} \ \{ \\
& \quad \bar{\alpha}_i := \bar{\alpha}_i - \rho_i y_i^\top \bar{q}; \\
& \quad \bar{\alpha}_i := \bar{\alpha}_i + \rho_i \alpha_i \bar{f}; \\
& \quad \bar{q} := \bar{q} + \bar{\alpha}_i s_i; \\
& \quad \bar{\alpha}_i := 0; \} \\
& \bar{g} := \bar{q}; \\
& \bar{q} := 0; \\
& \bar{f} := 0.
\end{aligned} \tag{3.5}$$

We now combine (3.4) and (3.5), omitting those instructions needed only for the evaluation of f (since we are only interested in $H_m g$). To save space in the adjoint portion of the code, we store all values of $\bar{\alpha}_i$ in the same location β and the values of \bar{q} in the location of q . After deleting all unnecessary instructions we obtain

$$\begin{aligned}
& q := g; \\
& \mathbf{for} \ i := m - 1 \ \mathbf{down\ to} \ 0 \ \mathbf{do} \ \{ \\
& \quad \alpha_i := s_i^\top q; \\
& \quad q := q - \rho_i \alpha_i y_i; \} \\
& q := H_0 q; \\
& \mathbf{for} \ i := 0 \ \mathbf{to} \ m - 1 \ \mathbf{do} \ \{ \\
& \quad \beta := \rho_i (\alpha_i - y_i^\top q); \\
& \quad q := q + \beta s_i; \}.
\end{aligned} \tag{3.6}$$

The value of $H_m g$ is placed in the vector q . Algorithm (3.6) is identical to the 2-loop formula used for the computation of the search direction in the L-BFGS method [6].

This derivation shows why algorithm (3.6) is efficient: it is based on the compact algorithm (3.4) computing f_m , and on the reverse mode of automatic differentiation, which is known to be very efficient.

4. References

- [1] Byrd, R.H., Nocedal, J. and Schnabel, R. (1992). Representations of quasi-Newton matrices and their use in limited memory methods, Tech. Rep. NAM-04, EECS Department, Northwestern University. 1
- [2] Gilbert, J.Ch., Le Vey, G., Masse, J. (1991). La différentiation automatique de fonctions représentées par des programmes. INRIA Research Report 1557. 2
- [3] Griewank, A. (1989). On automatic differentiation. In M. Iri, K. Tanabe, eds., *Mathematical Programming: Recent Developments and Applications*, pp. 83–108. Kluwer Academic Publishers. 2
- [4] Griewank, A., Corliss, G., eds. (1991). *Automatic Differentiation of Algorithms: Theory, Implementation, and Application*. Proceedings in Applied Mathematics 53. SIAM. 1

- [5] Juedes, D., “A taxonomy of automatic differentiation tools”, in *Automatic differentiation of algorithms: theory, implementation, and application*, Griewank, A. and Corliss, G. F. (editors), SIAM, Philadelphia, 1991. [2](#)
- [6] Nocedal, J. (1980). Updating quasi-Newton matrices with limited storage, *Mathematics of Computation*, Vol. 35, pp. 773–782. [1](#), [3](#), [4](#)
- [7] Rall, L.B. (1981). *Automatic Differentiation, Techniques and Applications*. Lecture Notes in Computer Science 120. Springer-Verlag, Berlin. [2](#)
- [8] Speelpenning, B. (1980). *Compiling fast partial derivatives of functions given by algorithms*. Ph.D. thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801. [2](#)