JOURNAL OF OPTIMIZATION THEORY AND APPLICATIONS: Vol. 37, No. 2, JUNE 1982

# Perspectives on Self-Scaling Variable Metric Algorithms[1]

S. S. OREN[2,3]

Communicated by H. Y. Huang

**Abstract.** Recent attempts to assess the performance of SSVM algorithms for unconstrained minimization problems differ in their evaluations from earlier assessments. Nevertheless, the new experiments confirm earlier observations that, on certain types of problems, the SSVM algorithms are far superior to other variable metric methods. This paper presents a critical review of these recent assessments and discusses some current interpretations advanced to explain the behavior of SSVM methods. The paper examines the new empirical results, in light of the original self-scaling theory, and introduces a new interpretation of these methods based on an $L$-function model of the objective function. This interpretation sheds new light on the performance characteristics of the SSVM methods, which contributes to the understanding of their behavior and helps in characterizing classes of problems which can benefit from the self-scaling approach.

**Key Words.** Unconstrained optimization, quasi-Newton methods, self-scaling variable metric methods, nonquadratic models.

## 1. Introduction

Self-scaling variable metric (SSVM) algorithms (Refs. 1–5) are a special class of variable metric methods for solving the unconstrained minimization problem:

$$\text{minimize } f(x), \qquad \text{where } x \in E^n \text{ and } f \in C^2.$$

[1] The subject of this paper was presented at the ORSA/TIMS National Meeting in New York, 1978.

[2] Associate Professor, Department of Engineering-Economic Systems, Stanford University, Stanford, California.

[3] This work was done while the author was with the Analysis Research Group, Xerox Palo Alto Research Center, Palo Alto, California.



0022-3239/82/0600-0137$03.00/0 

The general characteristic of such methods is that the successive approximations to the solution $\{x_k\}$ are obtained by a modified Newton iteration of the form

$$x_{k+1} = x_k - \alpha_k D_k g_k. \tag{1}$$

Here, $\alpha_k$ is a stepsize parameter,

$$g_k = \nabla f(x_k)^T,$$

and $D_k$ is an $n \times n$ symmetric matrix approximating the inverse Hessian $[\nabla^2 f(x_k)]^{-1}$, which is updated to satisfy the quasi–Newton condition:

$$D_{k+1} q_k = p_k, \tag{2}$$

where

$$q_k = g_{k+1} - g_k \quad \text{and} \quad p_k = x_{k+1} - x_k.$$

The updating is commonly done by a symmetric rank-two or rank-one modification of $D_k$, based on the discrepancy between the vector $p_k$ and its predicted value $D_k q_k$. Many popular updates such as the DFP and the BFGS formulas belong to a family known as Fletcher's (Ref. 6) convex class that can be expressed in the form

$$D_{k+1} = D_k - D_k q_k q_k^T D_k / q_k^T D_k q_k + \theta v_k v_k^T p_k p_k^T / p_k^T q_k, \tag{3}$$

where

$$v_k = (q_k^T D_k q_k)^{1/2} (p_k / p_k^T q_k - D_k q_k / q_k^T D_k q_k) \tag{4}$$

and $\theta \in [0, 1]$. This class has the property that, if $D_k$ is positive definite and $p_k^T q_k > 0$, then $D_{k+1}$ is also positive definite. Furthermore, if $f(x)$ is quadratic and

$$R_k = [\nabla^2 f(x)]^{1/2} D_k [\nabla^2 f(x)]^{1/2},$$

then the eigenvalues of $R_{k+1}$ are closer to unity than those of $R_k$.

The SSVM algorithms employ a modified version of the updates given by (3) and (4) in which the inverse Hessian approximation $D_k$ is scaled prior to its rank-two modification through multiplication by a scalar $\gamma_k$. This was originally motivated by the desire to decrease the condition number of the matrix $R_k$ which plays an important role in the convergence behavior of modified Newton algorithms [see Luenberger (Ref. 7)]. The scalar $\gamma_k$ is chosen so that the eigenvalues of the matrix $\gamma_k R_k$ are distributed around unity. Then, by moving the eigenvalues of $\gamma_k R_k$ closer to unity, the rank-two correction given by (3) and (4) will also decrease the ratio of the largest to smallest eigenvalue. This scheme is implemented by selecting $\gamma_k$

at each iteration such that

$$\gamma_k = (1-\varphi)(p_k^T q_k / q_k^T D_k q_k) + \varphi(p_k^T D_k^{-1} p_k / p_k^T q_k), \tag{5}$$

for some $\varphi \in [0, 1]$.

Since the original development of the SSVM algorithms, their practical merit has been debated. The main cause for the controversy is that scaling the inverse Hessian approximation beyond the first iteration relinquishes the property

$$D_n = [\nabla^2 f(x)]^{-1}$$

for a quadratic function. The above property has traditionally been a favored characteristic of variable metric algorithms, since it allows every $n$-step cycle of such methods to emulate an exact Newton step. This is compounded by the fact that, if the objective function is quadratic, then the purpose of the self-scaling approach can be accomplished by performing a self-scaling iteration only on the first step. This initial iteration does not have the adverse effect of subsequent SSVM iterations, since it acts merely as a device for selecting the initial approximation $D_0$, which is only restricted to be positive definite. Initial scaling can be sufficient even on nonquadratic problems, provided that the relative change in the norm of the Hessian over the relevant domain of the objective function is not too large. The need for scaling $D_k$ on each iteration arises only when the eigenvalues of $R_k$ drift rapidly away from unity, due to changes in the Hessian. The usual rank-two updates require an $n$-step cycle to bring all the eigenvalues of $R_k$ close to unity. Thus, if these eigenvalues drift away from unity relatively fast, or if the cycle is long, such updates cannot keep up with the change; hence, $D_k$ has to be scaled often to maintain the eigenvalues of $R_k$ distributed around unity. Such situations may occur when the objective function contains high powers of the variables or when the number of variables is large. Numerical experiments reported in Refs. 2, 5, 8, and 9 indeed confirm that the efficiency of SSVM increases in such situations.

## 2. Other Interpretations

Two recent papers by Brodlie (Ref. 10) and by Shanno and Phua (Ref. 11) present alternative interpretations and assess the merit of the self-scaling approach. Both papers interpret the scaling as an additional correction of the inverse Hessian approximation along the current search vector. Brodlie (Ref. 10) views the self-scaling mechanism as a device to get a better approximation of the curvature along the current search direction and attributes the advantages of SSVM methods to that feature. In light

of this interpretation, Brodlie compares the self-scaling approach to Biggs' (Ref. 12) method, which was indeed developed as a way to improve the approximation of the directional curvature. As a motivation for this interpretation, Brodlie points out that multiplying $D_k$ by

$$\gamma_k = p_k^T D_k^{-1} p_k / p_k^T q_k$$

ensures that the approximated curvature $p_k^T(\gamma_k D_k)^{-1} p_k$ will match the *true* curvature $p_k^T q_k$ prior to the rank-two correction. Shanno and Phua (Ref. 11), on the other hand, view the self-scaling mechanism as a device that improves the stepsize prediction and they attribute the advantages of SSVM methods to this characteristic. The latter interpretation is motivated by the observation, first given in Ref. 2, that, for a quadratic, the scaling factor

$$\gamma_k = p_k^T D_k^{-1} p_k / p_k^T q_k$$

equals the stepsize $\alpha_k$ that minimizes $f(x)$ along the line $x_k + \alpha D_k g_k$. Multiplying $D_k$ by $\gamma_k$ can, therefore, be interpreted as an adjustment of $D_k$, so that the minimizing stepsize $\alpha_k$ will be unity, as if the inverse Hessian approximation were exact. This interpretation suggests that, if exact line search is performed, the advantages of the SSVM algorithms will disappear.

Both papers present numerical results from which they conclude that SSVM methods are inferior to the BFGS (Ref. 13) algorithm. Shanno and Phua further show that the BFGS algorithm performs best when used with scaling on the first iteration. In both papers, however, the numerical results indicate that the SSVM methods outshine all other methods by a wide margin when applied to a quartic function of the form

$$f(x) = (x^T Q x)^2,$$

where $Q$ is a symmetric $n \times n$ nonsingular matrix. Brodlie attributes this observation to eratic behavior of the SSVM methods, while Shanno and Phua suggest that, on this type of function, the SSVM algorithms should be compared to Jacobson and Oksman's (Ref. 14) method which minimizes an $n$-dimensional homogenous function in $n+2$ steps. Neither of the interpretations given by the two papers, however, can satisfactorily explain the above phenomenon. Furthermore, none of the numerical results presented in the papers contradicts the original self-scaling theory. In fact, these results could have been predicted on the basis of that theory, since the only test problem in which the norm of the Hessian exhibits a significant relative change is the quartic function mentioned above. For this function, one can easily show that

$$4\lambda_{\max}\lambda_{\min}\|x\|^2 \leq \|\nabla^2 f(x)\| \leq 12\lambda_{\max}{}^2\|x\|^2,$$

where $\lambda_{\min}$ and $\lambda_{\min}$ are the smallest and largest eigenvalues of $Q$. Thus,

$$\|\nabla^2 f(x_{k+1})\|/\|\nabla^2 f(x_k)\| \leq 3(\lambda_{\max}/\lambda_{\min})\|x_{k+1}\|^2/\|x_k\|^2,$$

implying that, for any superlinearly converging algorithm, the ratio of the Hessian's norm on succesive interations approaches zero.

Finally, it should also be pointed out that the interpretations presented in the two aforementioned articles are limited to the particular scaling factor corresponding to $\varphi = 1$, and do not apply to scaling factors corresponding to other values of the parameter $\varphi$.

## 3. Linear $L$-Function Approach

Another interpretation of the self-scaling effect which takes a more global point of view has been proposed by Spedicato (Ref. 9). In his article, Spedicato develops the self-scaling update by considering an $L$-function model of the objective function, rather than a quadratic model. This model objective function has the form

$$f(x) = L\{\tfrac{1}{2}(x - x^*)^T Q(x - x^*) + c\}, \tag{6}$$

where $Q$ is a symmetric $n \times n$ positive definite matrix and $L(\cdot)$ is a strictly monotonically increasing, twice differentiable, real-valued function. Clearly, the function $f(x)$ given by (6) is minimized at the point $x^*$ which can be obtained by applying a variable metric algorithm with respect to the embedded quadratic

$$h(x) = \tfrac{1}{2}(x - x^*)^T Q(x - x^*) + c.$$

This approach, originated by Spedicato (Ref. 9), is implemented using the iteration (1) with $D_k$ approximating $[\nabla^2 h(x)]^{-1}$ and updated based on gradients of $h(x)$. Spedicato exploits the fact that the gradients $\nabla f(x)$ and $\nabla h(x)$ differ only by a scalar multiplier (depending on $x$), and develops a three-parameter updating formula for $D_k$ involving only gradients of $f(x)$. Spedicato also shows that, if $L(\cdot)$ is a simple linear scaling of its argument $h(x)$, then this three-parameter update reduces to the SSVM update. Unfortunately, for this special case, it is again sufficient to scale on the first iteration and proceed with a regular variable metric algorithm to achieve the desired effect of self-scaling. Thus, Spedicato's interpretation, as it stands does not explain the advantage of repeated self-scaling, which is most evident when $L(\cdot)$ is highly nonlinear.

## 4. New Interpretation

In this section, we introduce a new interpretation of the self-scaling methods based on a nonlinear $L$-function model of the objective function.

Let

$$f(x) = L\{h(x)\},$$

where $L(\cdot)$ is a scalar function. Then, a Newton iteration with respect to $h(x)$ can be expressed in the form

$$x_{k+1} = x_k[L_k'\nabla^2 h(x_k)]^{-1}\nabla f(x_k)^T, \tag{7}$$

where $L_k'$ denotes the first derivative of $L(\cdot)$. By analogy, the corresponding variable metric iteration with respect to $h(x)$ is given by (1), with $D_k$ approximating the matrix $[L_k'\nabla^2 h(x_k)]^{-1}$. One can easily show that

$$\nabla^2 f(x) = L''\nabla h(x)\nabla h(x)^T + L'\nabla^2 h(x), \tag{8}$$

where $L'$ and $L''$ are the first and second derivatives of $L(\cdot)$ evaluated at $h(x)$. Clearly, for

$$q_k = g_{k+1} - g_k \qquad \text{and} \qquad p_k = x_{k+1} - x_k,$$

we have

$$\nabla^2 f(x_{k+1})p_k \simeq q_k. \tag{9}$$

If, however, $x_{k+1}$ minimizes $f(x)$ along $p_k$, then

$$\nabla h(x_{k+1})p_k = 0.$$

Consequently, by (8) and (9),

$$L_{k+1}'\nabla^2 h(x_{k+1})p_k \simeq q_k. \tag{10}$$

This suggests that $D_k$ be updated to satisfy the usual quasi-Newton condition (2). The underlying philosophy of quasi-Newton methods is to satisfy the quasi-Newton condition by adding a rank-one or rank-two correction matrix to the best current estimate of the matrix to be approximated. As pointed out earlier, regular variable metric algorithms implement this philosophy by applying the correction to the matrix $D_k$ resulting from the previous update. This is consistent with the underlying quadratic model that assumes no change in the matrix approximated by $D_k$, namely, the inverse Hessian of the objective function. In our case, however, $D_k$ is the estimate of $[L_k'\nabla^2 h(x_k)]^{-1}$, while $D_{k+1}$ approximates $[L_{k+1}'\nabla^2 h(x_{k+1})]^{-1}$, which may be different. Thus, if $h(x)$ is assumed to be quadratic, then the proper estimate of $[L_{k+1}'\nabla^2 h(x_{k+1})]^{-1}$ prior to satisfying condition (2) is

$$D_k^* = \gamma_k D_k,$$

where

$$\gamma_k = L'_k / L'_{k+1}.$$

Consequently, the next approximation $D_{k+1}$ may be obtained by applying a rank-two correction to $D_k^*$, leading to an SSVM update.

One possible approach for implementing the updating procedure described above is to assume some specific function $L(\cdot)$ and obtain the corresponding expression for $\gamma_k$. Some alternative forms of $L(\cdot)$ that could be used for this purpose are given by Spedicato (Ref. 9). It should be noted, however, that if $L(\cdot)$ is a simple linear scaling of $h(x)$, then the above reasoning leads to setting

$$\gamma_k = 1, \qquad \text{for } k > 0.$$

This supports our earlier argument that a linearly scaled quadratic model of the objective function justifies only initial scaling.

An alternative approach for determining the scaling factor $\gamma_k$ is to estimate it from the information contained in the vectors $p_k$ and $q_k$. If $D_k$ and $\gamma_k$ were exact, i.e.,

$$D_k = [L'_k \nabla^2 h(x_k)]^{-1} \qquad \text{and} \qquad \gamma_k = L'_k / L'_{k+1},$$

then by (10) we would have

$$\gamma_k D_k q_k \simeq p_k.$$

It seems reasonable, therefore, to choose $\gamma_k$ to minimize the norm of the error between $\gamma_k D_k q_k$ and $p_k$ with respect to some metric $M_k$, i.e.,

$$\gamma_k = \arg\min_\gamma \|\gamma D_k q_k - p_k\|_{M_k}. \tag{11}$$

This leads to the equation

$$\gamma_k = q_k^T D_k M_k p_k / q_k^T D_k M_k D_k q_k. \tag{12}$$

Assuming that

$$q_k^T p_k > 0,$$

there exist a positive-definite matrix $S$ such that

$$S q_k = p_k.$$

Consequently, if $D_k$ is positive definite, then, for $\varphi \in [0, 1]$, the matrix $M_k^\varphi$, defined as

$$M_k^\varphi = (1-\varphi)(D_k^{-1} / q_k^T D_k q_k) + \varphi (D_k^{-1} S D_k^{-1} / p_k^T q_k), \tag{13}$$

is positive definite and satisfies

$$M_k^\varphi D_k q_k = (1-\varphi)(q_k / q_k^T D_k q_k) + \varphi (D_k^{-1} p_k / p_k^T q_k). \tag{14}$$

In particular, if $D_{k+1}$ is obtained by (3) and (4) (with $D_k$ replaced by $\gamma_k D_k$), then $D_{k+1}$ is positive definite and

$$D_{k+1}q_k = p_k,$$

so that the matrix $S$ in (13) may be replaced by $D_{k+1}$. Substituting (14) into (12) yields an equation for the scaling factor $\gamma_k$ identical to the original SSVM formula for $\gamma_k$ given by (5).

## 5. Discussion

The key concept underlying variable metric algorithms is the updating of the inverse Hessian approximation $D_k$. The correction of $D_k$ at each step is inferred from the discrepancy between a vector $p_k$ and its predicted value $D_k q_k$. The specific format of this correction depends on the underlying model of the objective function, which leads to the specific interpretation of the aforementioned discrepancy. Regular variable metric algorithms are based on a quadratic model of the objective function, thus treating the function as if the Hessian were constant during an $n$-step cycle. The discrepancy between $p_k$ and $D_k q_k$ is therefore attributed to the difference between $D_k$ and the true inverse Hessian along $p_k$, so the approximation $D_k$ is corrected by a rank-two or rank-one modification. SSVM algorithms, on the other hand, are based on an $L$-function model of the objective. Consequently, the discrepancy between $p_k$ and $D_k q_k$ is attributed to a change in scale as well as the difference between $D_k$ and the inverse Hessian along $p_k$. This leads to a two-step correction of $D_k$, utilizing the information contained in $p_k$ and $D_k q_k$ in two distinct ways. First, we infer the change in scale by finding the scalar multiplier $\gamma_k$ that will minimize the error between $p_k$ and $\gamma_k D_k q_k$, and then we apply a rank-two correction to compensate for the remaining discrepancy.

The above two approaches lead to different algorithms, the success of which depends largely on how well the underlying model of the objective function fits the objective function at hand. For small problems with relatively stable curvature, the quadratic model usually gives a better fit; and, consequently, algorithms based on that model perform better. On the other hand, in problems having a *flat bottom* (like the quartic function), an $L$-function model better represents the nature of the objective, thus giving an advantage to algorithms based on this model.

It should be pointed out that regular variable metric methods are capable of handling a certain degree of scale change. However, since their underlying model assumes no such change within an $n$-step cycle, it takes $n$-steps for such methods to accommodate a scale change. One would

expect, therefore, that, on large-scale problems, algorithms based on an $L$-function model will have a relative advantage, since such a model can better account for the change in scale that is likely to take place within a long $n$-step cycle. Nevertheless, even for large-scale problems, methods based on a quadratic model may do better if the scale is relatively stable.

To get the most out of an algorithm, one must understand its basic premises and the conditions under which it will perform best. Often, one has information about the problem at hand that enables him to determine *a priori* which model better represents the objective function, and choose the appropriate method. A classic example illustrating this point is the recent work by Dennis, Gay, and Welsch (Ref. 15). They have developed a new algorithm for solving the nonlinear least-squares problem. In this problem, one wishes to obtain a parameter vector $x \in E^n$ that minimizes the objective function

$$f(x) = \|r(x)\|^2,$$

where $r(x)$ is the $m$-dimensional residual vector, $m \geq n$, whose components

$$r_i(x) = y_i - \varphi_i(x)$$

denote the differences between the measurements $y_i$ and their corresponding predicted values $\varphi_i(x)$. The Hessian of this objective function is

$$\nabla^2 f(x) = \nabla r(x)^T \nabla r(x) + \sum_{i=1}^{m} r_i(x) \nabla^2 r_i(x). \tag{15}$$

The algorithm proposed in Ref. 15 is essentially a modified Newton algorithm for this problem, in which the first part of the Hessian, $\nabla r(x)^T \nabla r(x)$, is calculated analytically, while the second part is replaced by an approximation based on first derivatives. This approximation is updated on each iteration using a quasi-Newton approach. One of the outstanding difficulties for methods of this type arises when applied to zero-residual problems, i.e., problems for which the residuals can be reduced to zero. In such cases, as the residuals approach zero, the second term of the Hessian vanishes rapidly, and regular rank-one or rank-two updates of the approximation cannot catch up with this rapid change. To alleviate this difficulty, Dennis, Gay, and Welsch employed a self-scaling approach (referred to as *sizing*), which proved to be quite effective.

## 6. Conclusions

The results and discussion presented in this paper shed new light on and explain some of the observed empirical findings related to SSVM

algorithms. We hope that this will clarify some of the controversial issues related to these methods. The main objective of the paper, however, is to better characterize the situations in which SSVM performs best and facilitate the identification of special classes of problems that can benefit from this approach.

## References


1. OREN, S. S., *Self-Scaling Variable Metric Algorithms without Line Search for Unconstrained Minimization*, Mathematics of Computation, Vol. 27, pp. 873–885, 1973.
2. OREN, S. S., *Self-Scaling Variable Metric (SSVM) Algorithms, II, Implementation and Experiments*, Management Science, Vol. 20, pp. 863–874, 1974.
3. OREN, S. S., *On the Selection of Parameters in Self-Scaling Variable Metric Algorithms*, Mathematical Programming, Vol. 7, pp. 351–367, 1974.
4. OREN, S. S., and LUENBERGER, D. G., *Self-Scaling Variable Metric (SSVM) Algorithms, I, Criteria and Sufficient Conditions for Scaling a Class of Algorithms*, Management Science, Vol. 20, pp. 845–862, 1974.
5. OREN, S. S., and SPEDICATO, E., *Optimal Conditioning of Self-Scaling Variable Metric Algorithms*, Mathematical Programming, Vol. 10, pp. 70–90, 1976.
6. FLETCHER, R., *A New Approach to Variable Metric Algorithms*, Computer Journal, Vol. 13, pp. 317–322, 1970.
7. LUENBERGER, D. G., *Introduction to Linear and Nonlinear Programming*, Addison-Wesley Publishing Company, Reading, Massachusetts, 1973.
8. SPEDICATO, E., *Computational Experience with Quasi-Newton Algorithms for Minimization Problems of Moderately Large Size*, Report No. CISE-N-175, CISE Documentation Service, Segrate, Milano, Italy, 1975.
9. SPEDICATO, E., *A Variable Metric Method for Function Minimization Derived from Invariancy to Nonlinear Scaling*, Journal of Optimization Theory and Applications, Vol. 20, pp. 315–328, 1976.
10. BRODLIE, K. W., *An Assessment of Two Approaches to Variable Metric Methods*, Mathematical Programming, Vol. 12, pp. 344–355, 1977.
11. SHANNO, D. F., and PHUA, K. H., *Matrix Conditioning and Nonlinear Optimization*, Mathematical Programming, Vol. 14, pp. 149–160, 1978.
12. BIGGS, M. C., *Minimization Algorithms Making Use of Nonquadratic Properties of the Objective Function*, Journal of the Institute of Mathematics and Its Applications, Vol. 8, pp. 315–327, 1971.
13. BROYDEN, C. G., *The Convergence of a Class of Double Rank Minimization Algorithms, 2, The New Algorithm*, Journal of the Institute of Mathematics and Its Applications, Vol. 6, pp. 222–231, 1970.
14. JACOBSON, D. H., and OKSMAN, W., *An Algorithm that Minimizes Homogeneous Function of N Variables in N + 2 Iterations and Rapidly Minimizes*

*General Functions*, Journal of Mathematical Analysis and Applications, Vol. 38, pp. 535–552, 1970.

15. DENNIS, J. E., GAY, D. M., and WELSCH, R. E. *An Adaptive Nonlinear Least-Squares Algorithm*, Cornell University, Department of Computer Science, Technical Report No. 77-321, 1977.