Curvature in Conjugate Gradient Eigenvalue Computation with Applications to Materials and Chemistry Calculations

Alan Edelman* Tomás A. Arias† Steven T. Smith‡

Abstract

We illustrate the importance of using curvature information when performing constrained conjugate gradient minimization on manifolds and identify certain common and useful submanifolds in numerical analysis. This minimization approach promises to be useful in materials and chemistry calculations.

1 Introduction

The computation of the extreme eigenvalues of symmetric matrices may be posed as a constrained optimization problem. Minimization of $\frac{1}{2}\mathrm{tr}X^TAX$, where $X \in V_{n,k}$, the Stiefel manifold of $n \times k$ matrices with orthonormal columns, yields the lowest k eigenvectors of the matrix A. In this picture, non-linear eigenvalue problems become the simple generalization of the minimization of general smooth functions on the Stiefel manifold as well as the Grassman manifold of k dimensional subspaces of an n dimensional space.

The optimization approach for the eigenvalue problem has been considered only relatively recently. Commitment to this formulation raises the issue of whether conjugate gradient minimization is or can be an appropriate approach to solving the problem. In the numerical analysis community conjugate gradient style approaches to the solution of the symmetric eigenvalue problem through trace minimization may be found in work by Geradin, Sameh, and Wisniewski [7, 16]. Such techniques have also been used by Berry for singular value problems [3]. This work closely follows that of Smith [17] who shows how to solve such problems directly on the constraining manifold (so-called intrinsic techniques) by carefully taking into account the Riemannian curvature of the manifold.

A research area where conjugate gradient minimization of non-quadratic but smooth functions on the Stiefel manifold arises is the *ab initio* calculation of electronic structure within the local density approximation. Such approaches use only the charge and mass of electrons and atomic nuclei as input and have greatly furthered understanding of the thermodynamic properties of bulk materials [5], the structure and dynamics of surfaces [11, 13], the nature of point defects in crystals [14], and the diffusion and interaction of impurities in bulk materials [19]. Mirroring the numerical analysis community, optimization approaches have been only considered quite recently. Less than ten years ago, Car and Parrinello [6] in a watershed paper proposed minimization through simulated annealing. Teter and Gillan [8, 18] later introduced conjugate gradient based schemes

*Department of Mathematics Room 2-380, Massachusetts Institute of Technology, Cambridge, MA 02139, edelman@math.mit.edu.

†Department of Physics Massachusetts Institute of Technology, Cambridge, MA 02139, muchomas@mit.edu.

‡M.I.T. Lincoln Laboratory, 244 Wood Street; Room AN1-1K, Lexington, MA 02173-9108, stsmith@ll.mit.edu.

and demonstrated an order of magnitude increase in the convergence rate. These initial approaches, however, ignored entirely the effects of curvature on the choice of conjugate search directions. Taking the curvature into *partial* account using a generalization of the Riemannian projection led to a further improvement in computation times by over a factor of three under certain conditions [1], suggesting that significant gains can be made by accounting for the curvature *fully*.

The goal of this note is to bridge the gap between the various communities and to show how the geometrical framework developed by Smith [17] may be applied towards the development of practical algorithms for *ab initio* density functional computations.

2 Conjugate gradient eigenvalue computation

2.1 Conjugate gradient on the sphere and the importance of curvature

To understand the issues that arise in the conjugate gradient method on arbitrary Riemannian manifolds in [17], it is easiest to begin with the example of minimizing $\frac{1}{2}x^TAx$ on the sphere [17, Algorithm 4.5]. There are two important new features not seen in the classical algorithm: 1) straight line minimization is replaced by minimization along geodesics (great circles on spheres); correspondingly search directions and steepest descent directions are replaced with tangent vectors. A more subtle new feature is that 2) the Hessian must properly include the curvature terms associated with the manifold. Fortunately the latter may be done implicitly to sufficient accuracy.

On the sphere, tangent directions v at the point x satisfy $v^Tx = 0$. The steepest ascent direction of any smooth function that is defined on all of $\Re^n$ is the gradient of that function projected to the sphere. Therefore the function $f(x) = \frac{1}{2}x^TAx$, has steepest descent direction $-Ax + (x^TAx)x$. Of course, the eigenvectors of A have vanishing steepest descent direction, as they are stationary points of our function.

Minimization of $f(x) = \frac{1}{2}x^TAx$ along the geodesic that starts at x and heads in the direction v requires the solution of the 2 by 2 eigenvalue problem that is the projection of A on the plane spanned by x and v. Thus the exact minimization is roughly the solution of a quadratic equation. Of course if the function is more complicated, then some numerical optimization procedure is needed along the geodesic.

The final piece of the puzzle is the question of how to determine new search directions. The first requirement is that the new search direction is Hessian conjugate to the previous search direction. The second is that the new direction be in the space spanned by the steepest descent and the previous search directions. First of all the previous search direction must be rotated ("parallel transported" in the language of differential geometry) to the new position. Once we have the Hessian explicitly or implicitly, there is no difficulty in insuring conjugacy.

What is the correct Hessian for $\frac{1}{2}x^TAx$? The answer turns out not to be A, but rather $A - (x^TAx)I$. In Euclidean space the Hessian contains all the information necessary to compute the second derivative of the function along any straight line through a point x. On the sphere, the Hessian must contain the information necessary to compute the second derivative of the objective function along any geodesic through x. Since the geodesics are curved (great circles), there is a curvature term associated with the Hessian. (Curvature terms never show up when taking one derivative, but are essential when taking two.) A down-to-earth way to verify that $A - (x^TAx)I$ is the correct Hessian is to see if the second derivative in the direction v is $v^TAv - (x^TAx)$, when $\|v\| = 1$. To second order the geodesic may be written $p(t) = x + tv - (t^2/2)x$ since the acceleration must point inward towards the center. (To full order the geodesic is $p(t) = x\cos(t) + v\sin(t)$.) Then $(d^2/dt^2)p(t)^T Ap(t)$ is readily seen to be $v^TAv - x^TAx$. The term x^TAx exists because of the second order terms.

This Hessian has the nice property that it is invariant under shifts of the form $A \to A - sI$, while a Hessian that ignores the curvature term has no such nice property.

2.2 Conjugate Gradient Minimization is not Lanczos

There are well known links between the Lanczos algorithm to compute eigenvalues of a symmetric matrix and the classical conjugate gradient algorithm to minimize quadratic forms [9, p.494, 523]. However, it should be understood that the conjugate gradient method that we are using here to compute eigenvalues is not the same.

Our algorithm is similar to Lanczos in that it computes the extreme eigenvalues of symmetric matrices and is invariant under shifts $A \to A - sI$. In contrast to the classical conjugate gradient algorithm that solves linear systems by minimizing $\frac{1}{2}x^TAx - x^Tb$, this algorithm computes eigenvalues by minimizing x^TAx subject to constraints.

Classical conjugate gradient terminates in finitely many steps; of course no algorithm can generally compute eigenvalues exactly in a finite number of steps using the four elementary operations and square roots.

2.3 More than one eigenvalue

We now outline the important features of the generalization of the algorithm in Section 2.1 to the computation of the smallest k eigenvalues of an n by n matrix A. The function that we are now minimizing is $f(X) = \frac{1}{2}\mathrm{tr}(X^TAX)$, subject to the constraint that $X^TX = I_k$ These are $k(k+1)/2$ constraints on the point X thought of an element of nk dimensional Euclidean space.

The function $f(X) = \frac{1}{2}\mathrm{tr}(X^TAX)$ depends only on the vector space spanned by the columns of X, i.e. it is unchanged if X is replaced by XQ where Q is a $k \times k$ orthogonal matrix. Therefore f is a well defined function of the *Grassman manifold*, the set of k dimensional subspaces of an n dimensional space, which identifies all points in the Stiefel manifold related by an orthogonal transformation as a single point.

Let $I_{n,k}$ denote the first k columns of the identity matrix I_n. The tangent vectors V to the Stiefel manifold at $I_{n,k}$ are the n by k matrices V satisfying $V^TI_{n,k}$ is k by k anti-symmetric.

The curves $P(t) = \exp(tW)I_{n,k}$, where W is an n by n anti-symmetric matrix of the form

$$W = \begin{array}{c} \begin{array}{cc} k & n-k \end{array} \\ \begin{pmatrix} 0 & -Y^T \\ Y & 0 \end{pmatrix} \end{array}.$$

are all geodesics through $I_{n,k}$ on the Grassman manifold.

3 Application: Predicting material behavior from first principles

Our ability to compute *ab initio*, using only the charge and mass of electrons and atomic nuclei as input, the behavior of systems of everyday matter has advanced greatly in recent years. However, the computational demands of the approach and the attendant bounds on the size of systems which may be studied (roughly one hundred atoms) have limited the direct impact of the approach on materials and chemical engineering. Several *ab initio* applications which will benefit technology tremendously remain out of reach, requiring an order of magnitude increase in the size of addressable systems. Problems requiring the simultaneous study of thousands of atoms include defects in glasses (fiber optics communications), complexes of extended crystalline defects (materials' strength and processing), and large molecules (drug design).

The theoretical problem of interest is to find the smallest eigenvalue E_0 of the Schrödinger equation in the space of $3N$ dimensional antisymmetric functions,

$$H\psi = E_0\psi,$$

where the Hamiltonian operator H is defined by

$$H = \sum_{n=1\ldots N} [-\frac{1}{2}\nabla_n^2 + V_{ion}(r_n)] + \frac{1}{2}\sum_{1<n<m\leq N} \frac{1}{\|r_n - r_m\|^2}.$$

Here, N is the number of electrons in the system under study, now typically on the order of several hundred, r_i is the position of the ith electron, $V_{ion}(r)$ is the potential function due to the nuclei and inner electrons, and the second summation is recognized as the usual Coulomb interactions. Directly discretizing this equation at M grid-points in space would lead to absurdly huge eigenvalue problems where the matrix would be $M^N \times M^N$. This is not just a question of dense versus sparse methods, a direct approach is simply infeasible.

The fundamental theorems which make the *ab initio* approach tractable come from the density functional theory of Hohenberg and Kohn [10] and Kohn and Sham [12]. Density functional theory states that the ground states energy of a quantum mechanical system of interacting electrons and ions is equal to the solution of the problem of minimizing an energy function over all possible sets of N *three-dimensional* functions (electronic orbitals) obeying the constraints of orthonormality. Practical calculations generally use a finite basis to expand the orbitals, but for purposes of discussion, we may discretize the problem onto a finite spatial grid consisting of M points. The Kohn-Sham minimization then becomes,

$$\begin{aligned} E_0 &= \min_{X^TX=I_N} E(X) \\ &\equiv \min_{X^TX=I_N} \operatorname{tr}(X^THX) + f(\rho(X)), \end{aligned} \tag{1}$$

where each column of X is a different electronic orbital sampled on the spatial grid, ρ is the vector $\rho_i(X) \equiv \sum_n |X_{in}|^2$, H is an $M \times M$ matrix (single-particle Hamiltonian), and f is a function which we leave unspecified in this discussion. In full generality the X are complex, but the real case applies for physical systems of large extent that we envisage for this application [15], and we, accordingly, take X to be real in this note.

Recent advances in multicomputers have enabled such calculations on systems with several hundreds of atoms [2, 4]. Further improvements in memory and performance will soon make feasible computations with upwards of a thousand atoms. However, with growing interest in calculations involving larger systems has come the awareness that as the physical length of systems under study increases, the Hessian about the minimum of (1) becomes increasingly ill-conditioned and non-conjugate minimization approaches exhibit a critical slowing down [18]. This observation prompted workers [8, 18] to apply conjugate gradient concepts to problem. The promise of the geometric ideas in this note is that they may allow the effects of the curvature of the Stiefel manifold to be accounted properly for first time and thus significantly increase the efficiency of *ab initio* calculations.

The condition for the framework laid out in this note to be of practical use to the *ab initio* density-functional community is that inner product computation through the Hessian of $E(X)$ be no more computationally complex to evaluate than calculating the energy function $E(X)$ or maintaining the orthonormality constraints $X^T X = I_N$. A suitable form for this inner product computation is

$$
\begin{aligned}
(2) \qquad \frac{1}{2}\sum_{in,jm} Y_{in}\frac{\partial^2 E}{\partial X_{in}\partial X_{jm}} Z_{jm} &= \mathrm{tr}[Y^T(H+V)Z] + \sum_{ij}\sigma_i[2\frac{\partial^2 f}{\partial\rho_i\partial\rho_j}]\tau_j \\
&\quad -\mathrm{tr}[X^T(H+V)(XY^TZ)]
\end{aligned}
$$

where V is the diagonal matrix defined by $V_{ij} = (\partial f)/(\partial\rho_i)\delta_{ij}$, $\sigma_i \equiv \sum_n Y_{in}X_{in}$, $\tau_i \equiv \sum_n Z_{in}X_{in}$. Written this way, the first two terms of (2) have the same form and may be evaluated in the same manner as the corresponding terms in (1), with σ and τ playing roles similar to ρ. The third term, coming from the curvature, may be evaluated in the same way as the first term of (2) once given the object $[XY^TZ]$, which is no more computationally complex to obtain than the Gram-Schmidt orthonormalization of an object like X. Thus, geometrical conjugate gradient concepts bring no significant additional computational complexity to *ab initio* calculations.

The next section illustrates the gains afforded by taking account of the curvature in a simple test case. Work is currently underway to determine the increases in convergence rate associated with taking proper account of the curvature terms in full blown *ab initio* calculations. The potential utility of geometrically correct minimization algorithms to the electronic structure community is great as this community is nearing a qualitative change in the complexity and technological relevance of the problems it addresses.

4 Numerical Illustration

The figure to the right illustrates the computation of the smallest eigenvalue of a 500×500 matrix with eigenvalues 101 through 600 using three different algorithms: steepest descent, conjugate gradient without the curvature term in the Hessian, and conjugate gradient with the correct curvature term. Line minimizations along the geodesic were computed by solving the 2×2 eigenvalue problem, and each algorithm begins with the same random initial starting vector. We see that for this example, even the "wrong" approach to conjugate gradient is an improvement over steepest descent, but by properly taking into account curvature, we get further improvement.

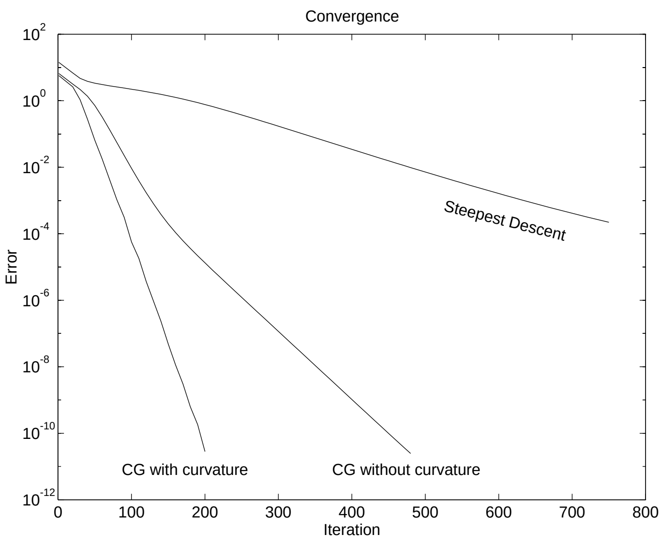

5 Acknowledgments

The first author is indebted to John Weare and his colleagues for an entire education in the local density approximation approach to first principles computations. Beth Ong also deserves special note for many conversations regarding the local density approximation and the Lanczos method.

References

[1] T.A. Arias, M.C. Payne and J.D. Joannopoulos, *Ab initio* molecular dynamics: analytically continued energy functionals and insights into iterative solutions, *Physical Review Letters 71*, (1992), 1077–1080.

[2] T.A. Arias and J.D. Joannopoulos, *Ab initio* theory of dislocation interactions: from close-range spontaneous annihilation to the long-range continuum limit, submitted to *Physical Review Letters*.

[3] M.W.Berry, Large-scale sparse singular value computations, *International J. Supercomputer Appl. 6* (1992), 13–49.

[4] K.D. Brommer, M. Needels, B.E. Larson, and J.D. Joannopoulos, *Ab initio* theory of the Si(111)-(7×7) surface reconstruction: a challenge for massively parallel computation, *Physical Review Letters 68*, (1992) 1355–1358.

[5] F. Buda, R. Car and M. Parrinello, Thermal expansion of *c*-Si via *ab initio* molecular dynamics. *Physical Review B 41*, (1990), 1680–1683.

[6] R. Car and M. Parrinello, Unified approach for molecular dynamics and density-functional theory, *Physical Review Letters 55*, (1985), 2471–2474.

[7] M. Geradin, The computational efficiency of a new minimization algorithm for eigenvalue analysis, *Journal Sound Vibration 19*, (1971), 319–331.

[8] M.J. Gillan, Calculation of the vacancy formation energy in aluminium, *Journal of Physics, Condensed Matter 1*, (1989) 689–711.

[9] G. Golub and C.F. Van Loan, *Matrix Computations, second edition*, Johns Hopkins University Press, Baltimore, 1989.

[10] P. Hohenberg and W. Kohn, Inhomogeneous Electron Gas, *Physical Review 136*, (1964), B864-B871.

[11] S. Ihara, S.L. Ho, T. Uda and M. Hirao, *Ab initio* molecular-dynamics study of defects on the reconstructed Si(001) surface, *Physical Review Letters 65*, (1990), 1909–1912.

[12] W. Kohn and L.J. Sham, Self-consistent equations including exchange and correlation effects, *Physical Review 140*, (1965), A1133-A1138.

[13] M. Needels, M.C. Payne, J.D. Joannopoulos, High order reconstructions of the Ge(100) surface, *Physical Review B 38*, (1988), 5543-5546.

[14] M. Needels, J.D. Joannopoulos, Y. Bar-Yam, and S.T. Pantelides, Oxygen complexes in Silicon, *Physical Review B 43*, (1991), 4208–4215.

[15] M.C.Payne,M.P.Teter,D.C.Allan,T.A.Arias,J.D.Joannopoulos, Iterative minimization techniques for *ab initio* total-energy calculations: molecular dynamics and conjugate gradients, *Rev. Mod. Phys. 64*, (1992), 1045–1097.

[16] A.H.Sameh and J.A.Wisniewski, A trace minimization algorithm for the generalized eigenvalue problem, *SIAM J. Num. Anal. 19* (1982), 1243–1259.

[17] S.T.Smith, *Geometric Optimization Methods for Adaptive Filtering*, PhD thesis, Department of Applied Mathematics, Harvard University, 1993.

[18] M.P.Teter, M.C.Payne, and D.C.Allan, Solution of Schrödinger's equation for large systems, *Physics Review B, 40* (1989), 255.

[19] C.G. Van de Walle, Y. Bar-Yam, S.T. Pantelides, *Theory of hydrogen diffusion and reactions in crystalline silicon*, *Physical Review Letters 60*, (1988) 2761–2764.