Manuscript submitted to
AIMS' Journals
Volume **X**, Number **0X**, XX **200X**

Website: http://AIMsciences.org

pp. **X–XX**

ORBITAL MINIMIZATION WITH LOCALIZATION

Weiguo Gao

School of Mathematical Sciences, Fudan University
Shanghai 200433, China

Weinan E

Department of Mathematics and Program in Applied and Computational Mathematics
Princeton University, Princeton, NJ 08544, USA

Dedicated to Professor Li Ta-Tsien on the occasion of his 70th birthday.

(Communicated by the associate editor name)

Abstract. Orbital minimization is among the most promising linear scaling algorithms for electronic structure calculation. However, to achieve linear scaling, one has to truncate the support of the orbitals and this introduces many problems, the most important of which is the occurrence of numerous local minima. In this paper, we introduce a simple modification of the orbital minimization method, by adding a localization step into the algorithm. This localization step selects the most localized representation of the subspace spanned by the orbitals obtained during the intermediate stages of the iteration process. We show that the addition of the localization step substantially reduces the chances that the iterations get trapped at local minima.

1. **Introduction.** Orbital minimization (OM, in short) is one of the most promising linear scaling algorithms for computing the electronic structure of materials or molecules. In its simplest form, the problem is to minimize the functional:

$$E(\Psi) = \mathrm{Tr}\left((\Psi^T\Psi)^{-1}\Psi^T H\Psi\right). \tag{1}$$

Here H is the Hamiltonian operator acting on R^N, $\Psi = (\psi_1, \cdots, \psi_M)$ is the $N \times M$ matrix formed by the orbitals $\{\psi_1, \cdots, \psi_M\}$. It is easy to see that the functional E is invariant under any non-singular linear transformations:

$$E(\Psi) = E(\Psi G) \tag{2}$$

for any non-singular $M \times M$ matrix G. Hence E should really be viewed as a functional defined on the Grassmann manifold that consists of the M dimensional subspaces of R^N. In particular, the minimizer of E should be a subspace of R^N. Indeed it is easy to check that the minimizer of E is simply the eigen-subspace of H corresponding to its M smallest eigenvalues. In fact, we can view E as the generalization of the Rayleigh quotient for subspaces.

2000 *Mathematics Subject Classification.* Primary: 46N40; Secondary: 35Q40.

Key words and phrases. Orbital minimization, localization, conjugate gradient.

This work is supported in part by ONR grant N00014-01-1-0674, DOE grant DOE DE-FG02-03ER25587, NSF grants DMS-0708026 and DMR-0611562, with additional support for Gao by the Special Funds for Major State Basic Research Projects of China, grant 2005CB321700 and the 111 Project.

From another angle, we see that the minimizer of E is not unique: If Ψ is a minimizer, then any non-singular linear transformation of Ψ is also a minimizer, since it spans the same subspace. Therefore it is not surprising that choosing the appropriate representation of the subspace is a significant aspect of this problem. This was the motivation of [2] and is also the focus of the present paper.

From a numerical viewpoint, this problem can be solved using more or less standard optimization techniques such as the steepest decent and the conjugate gradient method, as illustrated below. This is the basis of OM. Different versions of OM have been proposed since the 1990s, see [4, 7] and the references therein. Besides (1), another popular form of OM is:

$$\tilde{E}(\Psi) = \mathrm{Tr}\left((2I - \Psi^T\Psi)\Psi^T H\Psi\right). \tag{3}$$

As was observed in [14, 15], if H is negative definite, the inverse of the overlap matrix $\Psi^T\Psi$ in (1) can be replaced by $2I-\Psi^T\Psi$ without changing the global minimum. By shifting H, i.e. replacing H by $H - \mu I$, one can always make H negative definite, and this does not change the solution of the problem. Therefore, minimizing E is equivalent to minimizing $\tilde{E}$. Compared with (1), (3) has the advantage that it does not require calculating the inverse of the overlap matrix. Furthermore, one can show that under fairly general conditions, orthogonality between the orbitals is guaranteed to hold for the global minimizer [13, 16]. In fact, (3) can also be derived by considering the constrained minimization problem:

$$E(\Psi) = \mathrm{Tr}\left(\Psi^T H\Psi - (\Psi^T\Psi - I)\Lambda\right) \tag{4}$$

where Λ is an $M \times M$ matrix consisting of the Lagrange multipliers for the constraints:

$$\Psi^T\Psi = I. \tag{5}$$

For this reason, it is natural to associate (3) or (4) with orthogonal orbital minimization and (1) with non-orthogonal orbital minimization. Further references on OM can be found in [1, 5, 9, 19, 20, 21]. For the comparison of different schemes, we refer to [11, 17].

In this paper, we will focus on the nonorthogonal representation (1) only. There are two main reasons. The first is that nonorthogonal orbitals can be more localized than orthogonal ones. As we will see below, the main point of the new OM algorithm is to add a localization step. The performace of the new algorithm critically depends on how localized the orbitals are. The second is that the nonorthogonal formulation (1) does not contain other numerical parameters such as the μ orthogonal formulation. As has been observed before [17], the rate of convergence depends quite sensitively on the value of μ. In this regard, the performance of the nonorthogonal formulation is more stable.

In practice, in order to obtain linear scaling, one has to truncate the support of the orbitals. Let L be the size of the support of the orbitals after truncation. We use the notation

$$\Psi_L = ((\psi_1)_L, \cdots, (\psi_M)_L)$$

to denote the truncated orbitals, see Figure 1 as an example. The truncated functional takes the form:

$$E_L(\Psi_L) = \mathrm{Tr}\left((\Psi_L^T\Psi_L)^{-1}\Psi_L^T H\Psi_L\right). \tag{6}$$

We can view this as the restriction of the original energy functional to the space of truncated orbitals Ψ_L.

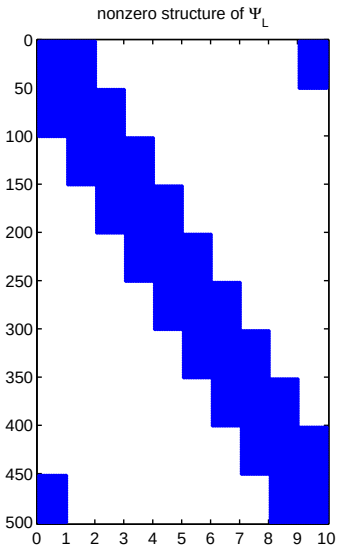

FIGURE 1. The support of the truncated orbitals. $N = 500$, $M = 10$, and the size of the support for each orbital $L = 150$.

The intuitive reason for the truncation of the orbitals is that orbitals are often localized [10], i.e. each orbital decays very fast away from its center. In particular, for insulators, it is generally believed and to some extend proved that this decay can be made at an exponential rate [10, 18]. Therefore, in principle we should only incur a small error if we truncate the support of the orbitals beyond certain size.

However, it is widely observed that the performance of this truncated version of OM is much less stable than the original OM. The most serious difficulty is that the truncated problem has numerous local minima and the numerical solutions are often trapped in these local minima. To illustrate this point, we show in Figure 2 the results of OM for a simple model problem to be discussed later (the first example in Section 3). We chose $N = 500, M = 10, L = 150$. Both standard steepest decent (Algorithm 1 in Section 3) and conjugate gradient (Algorithm 3 in Section 3) algorithms are used as the optimization procedure. We initialize the algorithms with random initial conditions, i.e. each component of the vectors are independent random variables uniformly distributed in $[0, 1]$. We stop the iteration when the relative change of energy is less than some tolerance ϵ and we used a variety of values of ϵ, ranging from 10^{-7} to 10^{-10}. The results are rather insensitive to the value of ϵ used.

Figure 2 is a plot of the energy of the converged solutions against an index that numbers the initial conditions. It is obvious from Figure 2 that for a substantial percentage of the initial conditions, the numerical solutions are trapped at the local minima (or saddle point). This is a rather unpleasant feature of the algorithm. As expected, the results are improved as we increase the size of L (see Figure 3).

At an intuitive level, it is easy to see that something like this might happen. The original problem is invariant under any non-singular linear transformation of the orbitals, see (2). However, after truncation (6), this invariance property is lost. Instead, we are looking for solutions in the subspace spanned by the truncated

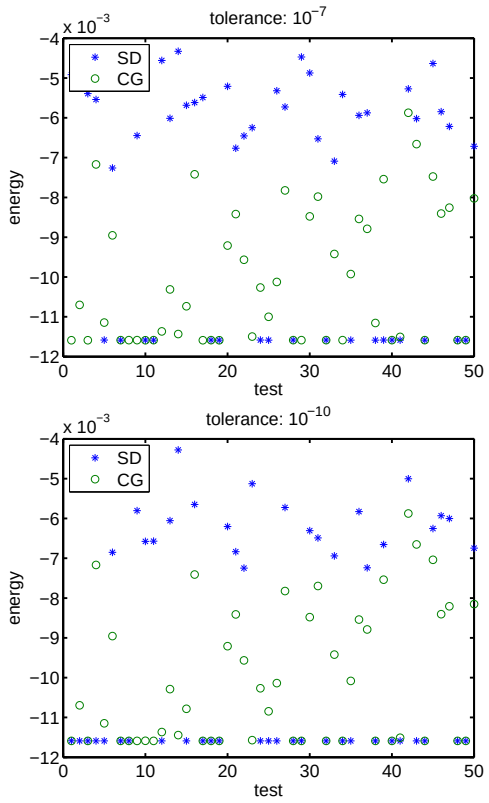

FIGURE 2. 50 random tests of the steepest descent and conjugate gradient algorithms for the second order discretization of the operator $\mathcal{H} = -\frac{1}{2}\nabla^2 + V(x)$ in Section 3, $N = 500$, $M = 10$, $L = 150$, the error tolerance ϵ is chosen to be 10^{-7} and 10^{-10} respectively.

vectors. Therefore, it is entirely possible that some of the saddle points of the original problem will generate nearby local minima for the new problem.

Our strategy for overcoming this difficulty is to look for representations of the subspaces that are closest to the truncated space. This is done by introducing a localization step into the algorithm. This localization step substantially reduces the chances that the numerical solution gets trapped at the local minima. In fact, also shown in Figure 1 are the results of the OM with localization (hereafter denoted as OM_L). It is quite obvious that there are a lot fewer incidences of trapping by the local minima.

The idea of adding a localization step was first introduced in the localized subspace iteration (LSI) algorithm in [2]. [2] also analyzed systematically the performance of different localization procedures, and in particular the influence of the choice of the weight function for weight-function-based localization. However, it should be noted that the localization procedure itself has been quite popular for some time, particularly following the seminal work by Marzari and Vanderbilt [12].

We will first present a localization procedure which appears to be simpler than the standard ones based on weight functions.

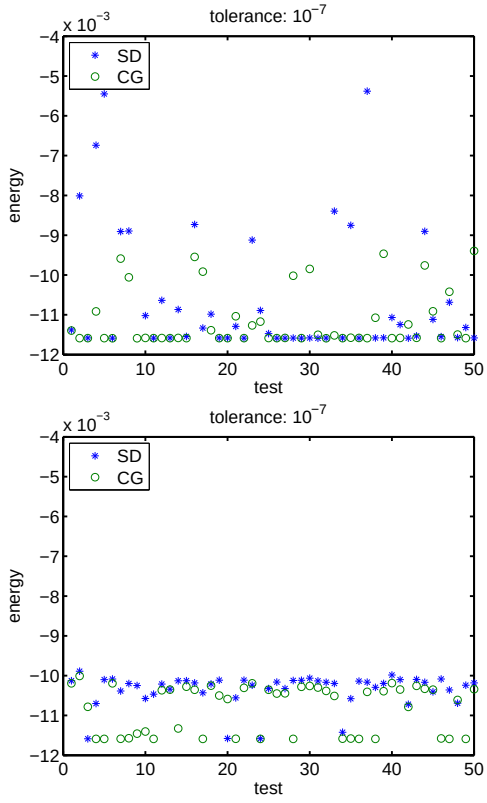

FIGURE 3. 50 random tests of the steepest descent and conjugate gradient algorithms for the second order discretization of the operator $\mathcal{H} = -\frac{1}{2}\nabla^2 + V(x)$ in Section 3, $N = 500$, $M = 10$, $\epsilon = 10^{-7}$, L is set to be 250 and 350 respectively.

2. **The localization procedure.** In this section, we discuss the localization procedure that we will use. Given a set of orbitals that composes the matrix Ψ, let V_Ψ be the subspace spanned by these orbitals, which are the column vectors of Ψ. The problem can be formulated as follows: Given Ψ, we would like to find a set of truncaed orbitals, denoted by Ψ_0 such that V_{Ψ_0} gives the best approximaiton of V_Ψ.

To give a precise formulation of this problem, let us first clarify the notion of truncated orbitals. Assume for the moment that the Hamiltonian H is the discretization of a differential operator on a mesh with N grid points. Let us select M representative points out of the N grid points. For example, we may select M equally-spaced grid points. Denote them by $x_i, i = 1, \cdots, M$. Let S_i be a neighborhood of x_i. $\Psi_0 = (\psi_1^0, \cdots, \psi_M^0)$ is a set of localized orbitals if ψ_i^0 vanishes outside S_i, for $i = 1, \cdots, M$ (see Figure 1). If $\Psi = (\psi_1, \cdots, \psi_M)$ is an arbitrary set of orbitals, we will denote by $\Psi_L = ((\psi_1)_L, \cdots, (\psi_M)_L)$ the set of orbitals obtained by simply truncating ψ_i outside S_i for each i.

We can now formulate the problem as follows: We would like to find a $M \times M$ non-singular matrix G such that the new representation ΨG of V_Ψ minimizes the cutoff $\Psi G - (\Psi G)_L$ in some sense. For computational efficiency, we will use the Frobenius norm to measure the error due to cutoff.

To proceed further, let us observe that minimizing the functional $\|\Psi G-(\Psi G)_L\|_F$ can be done column by column, i.e. if we write $G = (\alpha_1, \ldots, \alpha_M)$ in terms of the column vectors, we obtain M minimization problems

$$\min \|\Psi\alpha_i - (\Psi\alpha_i)_L\|_2, \quad i = 1, \ldots, M. \tag{7}$$

These problems can be solved independently.

It is obvious that in (7) the components of $\Psi\alpha_i$ and $(\Psi\alpha_i)_L$ are equal inside the localization region S_i for each i. Therefore we only need to consider the remaining components not included in the localization region, which will be truncated. We will use $\Psi^{(i)}$ to represent the resulted orbitals after deleting the rows corresponding to the grid points in the localization region. Then the problem becomes to minimize

$$\|\Psi^{(i)}\alpha_i\|_2, i = 1, \ldots, M$$

with some constraint on α_i.

One possible choice of the constraint is $\|\alpha_i\|_2 = 1$. We then obtain the following minimization problem

$$\min_{\|\alpha_i\|_2=1} \|\Psi^{(i)}\alpha_i\|_2 \tag{8}$$

This is the problem of finding the smallest singular value for $\Psi^{(i)}$ and α_i is the corresponding right singular vector. Another natural constraint is $\alpha_i^T e = 1$ where $e = (1, \cdots, 1)^T$. The minimization problem

$$\min_{\alpha_i^T e=1} \|\Psi^{(i)}\alpha_i\|_2 \tag{9}$$

is explicitly solved by

$$\alpha_i = \frac{\beta_i}{\beta_i^T e}$$

where $\beta_i = ((\Psi^{(i)})^T\Psi^{(i)})^{-1}e$. In practice, we do not have to compute $(\Psi^{(i)})^T\Psi^{(i)}$. A stable QR algorithm [8] can be used to solve the constrained least-square problem (9). Since typically $M << N$, the cost of the localization step in each iteration is acceptable, typically less than 10% in our runs.

If we think about the changes of the support of the orbitals, we get the following picture: After truncation, the support of the i-th orbital becomes S_i. The support is then increased, due to the application of the Hamiltonian operator on the orbitals, which is a typical step in many optimization algorithms. Since in the cases we are interested in, the Hamiltonian operator is very local, therefore the application of the Hamiltonian operator only increases the support of the orbitals slightly, as illustrated in Figure 4. α_i has contributions only from the columns that correspond to the components shown in red. $\Psi^{(i)}$ is obtained by removing the rows corresponding to the components shown in blue. We can even think of $\Psi^{(i+1)}$ as a low-rank correction of $\Psi^{(i)}$ since $\Psi^{(i+1)}$ can be obtained from $\Psi^{(i)}$ by removing and adding a few rows and columns. In this way, we can use both low-rank update versions of the SVD and QR algorithms [8] to get the solution to (7). The support goes back to S_i after the next truncation step.

To get an idea of how the localization step changes the performance of the algorithm, we show in Figure 5 the error due to truncation with (right) or without (left) the localization step added. Obviously the truncated components are much smaller with the localization step.

The localization procedure proposed here appears to be simpler than the one proposed in [2]. For example, $\Psi^{(i)}$ does not include the rows in the localization region

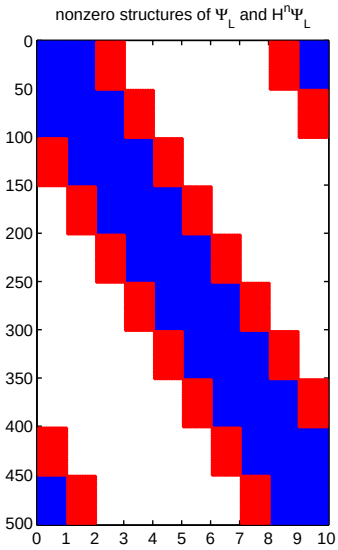

FIGURE 4. Expansion of the support of the orbitals applying the Hamiltonian operator 50 times, $N = 500$, $M = 10$, $L = 150$.

and this leads to smaller problems. In addition, we avoid solving the generalized eigenvalue problem.

3. **OM with localization.** In this section, we will show how to incorporate the localization step into the different variants of the OM algorithm. We will test these algorithms on two model problems. The first is an example for which H is a second order discretization of the operator $\mathcal{H} = -\frac{1}{2}\nabla^2 + V(x)$, where

$$V(x) = \alpha \sum_j e^{-(x-r_j)^2/(2\sigma^2 d^2)}$$

is the external potential. The system is defined on the the interval $[0, 10]$ with periodic boundary condition. We used the values $r_j = 0.5, 1.5, \ldots, 9.5$, $\alpha = -100.0, d = 1.0$ and $\sigma = 0.01$. The discretization parameters N, M and L are chosen as $N = 500$, $M = 10$ and $L = 150, 250, 350$ respectively. Here L is the size of the support S_i (number of grid points in S_i).

The parameters have been chosen to correspond to the case of an insulator, i.e. there is a gap between the eigenvalues corresponding to the eigen-subspace spanned by the optimal orbitals and the rest of the spectrum. In the context of electronic structure, in this case, one can find exponentially decaying orbitals [10]. Therefore, we expect the localization procedure to work quite well, and if we find the right set of representative orbitals, the error due to truncation can be made quite small.

The second example is for the case when H is simply a second order discretization of the one-dimensional Laplacian operator, which is equivalent to setting $\alpha = 0$ in the first example. Even though the operator is very simple, this is actually a tough example since there are no spectral gap for this problem, i.e. it corresponds to the case of a metal in the setting of electronic structure analysis. In this case, we expect the orbitals to decay at a slower rate. Therefore, the error caused by truncation should be larger compared with the previous case.

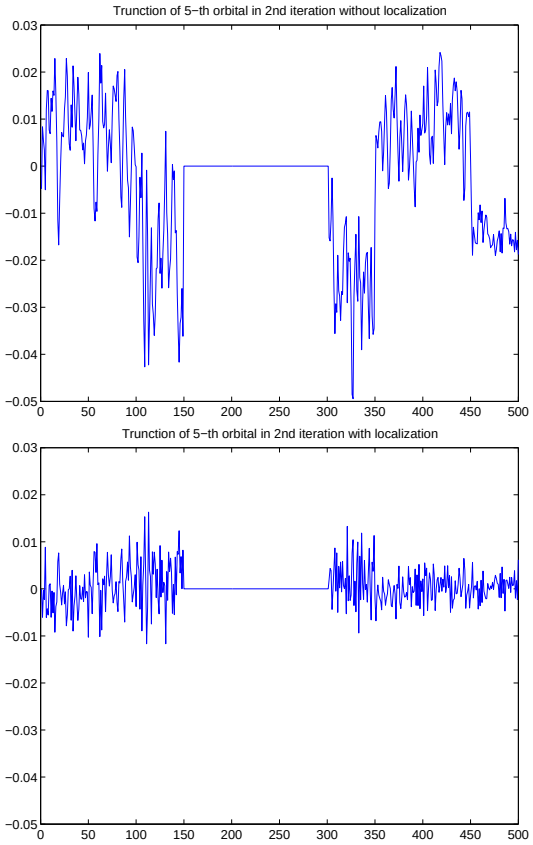

FIGURE 5. Truncated components of the fifth orbital in the second iteration of the steepest descent algorithm. The result of direct truncation is shown on the left. The result of truncation after localization is shown on the right.

3.1. **Steepest decent algorithm.** We start with the steepest decent algorithm. It is easy to verify that the gradient of the truncated functional $E_L(\Psi_L)$ is the truncation of the gradient of the functional $E(\Psi_L)$ [13]

$$\nabla E_L(\Psi_L) = (\nabla E(\Psi_L))_L.$$

A direct application of the steepest decent algorithm with truncation is as follows:

Algorithm 1 Steepest Decent with Truncation (SD)
Input: Ψ_L^0, initial guess of Ψ_L
1) Start: $D_L^0 = -\nabla E_L(\Psi_L^0)$
2) Iterate: For $k = 0, 1, \ldots,$
3) $\alpha_k = \text{argmin}_\alpha E_L(\Psi_L^k + \alpha D_L^k)$
4) $\Psi_L^{k+1} = \Psi_L^k + \alpha_k D_L^k$
5) $D_L^{k+1} = -\nabla E_L(\Psi_L^{k+1})$
6) Check convergence
7) EndDo
Output: Ψ_L^k, approximation of $(\Psi_L)^*$ (the minimizer of E_L).

To add the localization step, at the kth iteration, we replace D_L^k by the gradient D^k of the original functional (1). The new orbitals Ψ^{k+1} are then calculated from D^k. However, instead of direct truncation, we first use the localization procedure to find a nonsingular transformation G described in the last section. This leads to the following algorithm:

Algorithm 2 Steepest Decent with Localization (SD_L)
Input: Ψ_L^0, initial guess of Ψ_L
1) Start: $D^0 = -\nabla E(\Psi^0)$
2) Iterate: For $k = 0, 1, \ldots,$
3) $\quad \alpha_k = \mathrm{argmin}_\alpha E(\Psi_L^k + \alpha D^k)$
4) $\quad \Psi^{k+1} = \Psi_L^k + \alpha_k D^k$
5) $\quad G = \mathrm{argmin}_G \|\Psi^{k+1} G - (\Psi^{k+1} G)_L\|_F$
6) $\quad \Psi_L^{k+1} = (\Psi^{k+1} G)_L$
7) $\quad D^{k+1} = -\nabla E(\Psi_L^{k+1})$
8) $\quad$ Check convergence
9) EndDo
Output: Ψ_L^k, approximation of $(\Psi_L)^*$.

The results for the first example are shown in Figure 6. One can see clearly that

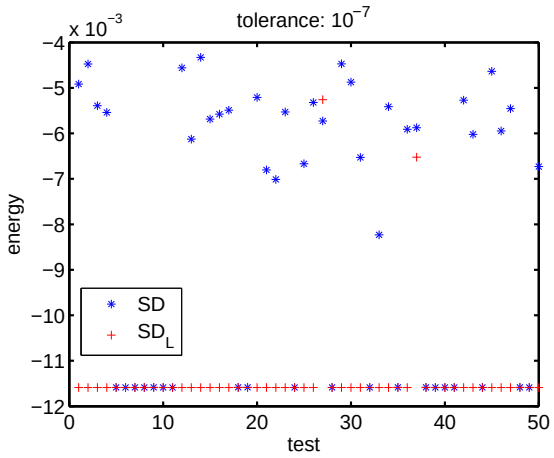

FIGURE 6. 50 random tests of the steepest descent (SD) and steepest decent with localization (SD_L) algorithms for the second order discretization of the operator $\mathcal{H} = -\frac{1}{2}\nabla^2 + V(x)$, $N = 500$, $M = 10$, $\epsilon = 10^{-7}$, $L = 150$.

results from Algorithm 2 show much less trapping at local minima than results from Algorithm 1.

3.2. **Conjugate gradient method.** If we use the conjugate gradient method with direction truncation of the orbitals, we obtain the following algorithm:

Algorithm 3 Conjugate Gradient with Truncation (CG)
Input: Ψ_L^0, initial guess of Ψ_L
1) Start: $R_L^0 = -\nabla E_L(\Psi_L^0), D_L^0 = R_L^0$
2) Iterate: For $k = 0, 1, \ldots,$
3) $\quad \alpha_k = \mathrm{argmin}_\alpha E_L(\Psi_L^k + \alpha D_L^k)$

4) $\Psi_L^{k+1} = \Psi_L^k + \alpha_k D_L^k$
5) $R_L^{k+1} = -\nabla E_L(\Psi_L^{k+1})$
6) $\beta_{k+1} = \dfrac{< R_L^{k+1}, R_L^{k+1} >}{< R_L^k, R_L^k >}$
7) $D_L^{k+1} = R_L^{k+1} + \beta_{k+1} D_L^k$
8) Check convergence
9) EndDo

Output: Ψ_L^k, approximation of $(\Psi_L)^*$.

When adding the localization step, we also need to transform the search direction accordingly.

Algorithm 4 Conjugate Gradient with Localization (CG_L)
Input: Ψ_L^0, initial guess of Ψ_L
1) Start: $R^0 = -\nabla E(\Psi^0), D^0 = R^0$
2) Iterate: For $k = 0, 1, \ldots,$
3) $\alpha_k = \mathrm{argmin}_\alpha E(\Psi_L^k + \alpha D^k)$
4) $\Psi^{k+1} = \Psi_L^k + \alpha_k D^k$
5) $G = \mathrm{argmin}_G \|\Psi^{k+1} G - (\Psi^{k+1} G)_L\|_F$
6) $\Psi_L^{k+1} = (\Psi^{k+1} G)_L$
7) $R^{k+1} = -\nabla E(\Psi_L^{k+1})$
8) $\beta_{k+1} = \dfrac{< R^{k+1}, R^{k+1} >}{< R^k, R^k >}$
9) $D^{k+1} = R^{k+1} + \beta_{k+1} (D^k G)_L$
10) Check convergence
11) EndDo

Output: Ψ_L^k, approximation of $(\Psi_L)^*$.

Figure 7 shows the numerical results of conjugate gradient method for the same example. We see that the new algorithm with the localization step added performs

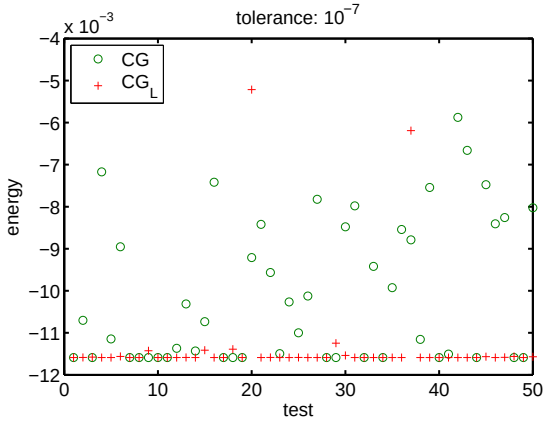

FIGURE 7. 50 random tests of the conjugate gradient (CG) and conjugate gradient with localization (CG_L) algorithms for the second order discretization of the operator $\mathcal{H} = -\frac{1}{2}\nabla^2 + V(x)$, $N = 500$, $M = 10$, $\epsilon = 10^{-7}$, $L = 150$.

much better.

3.3. **Grassmann conjugate gradient method.** As we remarked earlier, the reason for the generation of numerous local minima in the truncated functional (or functional defined on truncated orbitals) is that the space of the truncated orbitals is no longer invariant under non-singular linear transformations. The original functional is well-defined on the Grassmann manifold. This is no longer the case for the truncated functional. To compensate for this, one might use a version of the conjugate gradient method that respects the invariance under nonsingular linear transformation. This is the Grassmann conjugate gradient method, which is also naturally defined on the Grassmann manifold of R^N [3]. In this method, the gradient must lie in the tangent plane of the Grassmann manifold.

If we use this version of the conjugate gradient method, we obtain the following orbital minimization algorithms, with or without localization.

Algorithm 5 Grassmann Conjugate Gradient with Truncation (GCG)
Input: Ψ_L^0, initial guess of Ψ_L
1) Start: $R_L^0 = -\nabla E_L(\Psi_L^0), D_L^0 = R_L^0$
2) Iterate: For $k = 0, 1, \ldots,$
3) $\alpha_k = \text{argmin}_\alpha E_L(\Psi_L^k + \alpha D_L^k)$
4) $D_{new}^k = D_L^k - \alpha_k \Psi_L^k ((\Psi_L^k)^T \Psi_L^k)^{-1} (D_L^k)^T D_L^k$
5) $R_{new}^k = R_L^k - \alpha_k \Psi_L^k ((\Psi_L^k)^T \Psi_L^k)^{-1} (D_L^k)^T R_L^k$
6) $\Psi_L^{k+1} = \Psi_L^k + \alpha_k D_L^k$
7) $R_L^{k+1} = -\nabla E_L(\Psi_L^{k+1})$
8) $\beta_{k+1} = \dfrac{< R_L^{k+1}, R_L^{k+1} >}{< R_{new}^k, R_{new}^k >}$
9) $D_L^{k+1} = R_L^{k+1} + \beta_{k+1} D_{new}^k$
10) Check convergence
11) EndDo
Output: Ψ_L^k, approximation of $(\Psi_L)^*$.

Algorithm 6 Grassmann Conjugate Gradient with Localization (GCG_L)
1) Start: $R^0 = -\nabla E(\Psi^0), D^0 = R^0$
2) Iterate: For $k = 0, 1, \ldots,$
3) $\alpha_k = \text{argmin}_\alpha E(\Psi_L^k + \alpha D^k)$
4) $D_{new}^k = D_L^k - \alpha_k \Psi_k ((\Psi_L^k)^T \Psi_L^k)^{-1} (D_L^k)^T D_L^k$
5) $R_{new}^k = R_L^k - \alpha_k \Psi_k ((\Psi_L^k)^T \Psi_L^k)^{-1} (D_L^k)^T R_L^k$
6) $\Psi^{k+1} = \Psi_L^k + \alpha_k D^k$
7) $G = \text{argmin}_G \|\Psi^{k+1} G - (\Psi^{k+1} G)_L\|_F$
8) $\Psi_L^{k+1} = (\Psi^{k+1} G)_L$
9) $R^{k+1} = -\nabla E(\Psi_L^{k+1})$
10) $\beta_{k+1} = \dfrac{< R^{k+1}, R^{k+1} >}{< R_{new}^k, R_{new}^k >}$
11) $D^{k+1} = R^{k+1} + \beta_{k+1} (D_{new}^k G)_L$
12) Check convergence
13) EndDo
Output: Ψ_L^k, approximation of $(\Psi_L)^*$.

The results are shown in Figure 8.

4. **An example without spectral gap.** We now consider the second example, i.e., one-dimensional Laplacian operator. For this example, we chosethe system size to be $N = 100$. The results are shown in Figures 9, 10 and 11 respectively.

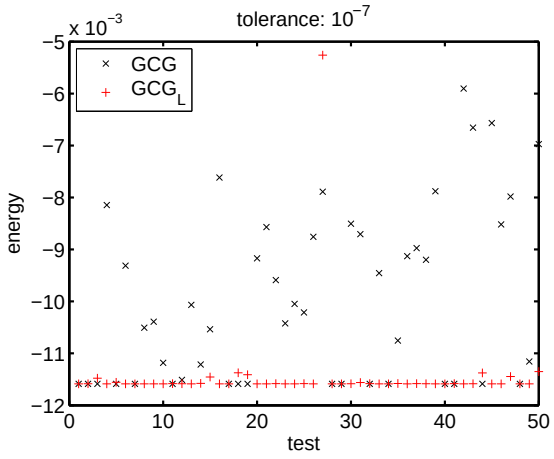

FIGURE 8. 50 random tests of Grassmann conjugate gradient (GCG) and Grassmann conjugate gradient with localization (GCG_L) algorithms for the second order discretization of the operator $\mathcal{H} = -\frac{1}{2}\nabla^2 + V(x)$, $N = 500$, $M = 10$, $\epsilon = 10^{-7}$, $L = 150$.

The condition number of the first example is smaller than that of the second, since there is a spectral gap for the first example [6, 10]. Consequently, it requires more iterations for convergence for the second example, and this is what we observed.

On the other hand, the algorithms with the localization step always run faster and give more accurate results for both examples. To demonstrate the of the localization step, we increase the size of system N to 4000, the number of orbitals $M = 20$, σ to -1000.0. The performance of conjugate gradient with and without localization is presented in Table 1. We see that with the localization step, the conjugate gradient

Test	1	2	3	4	5
CG_L	133	155	170	158	148
CG	201	325	359	315	164

TABLE 1. Numbers of iterations of 5 random tests of the conjugate gradient (CG) and conjugate gradient with localization (CG_L) algorithms for the second order discretization of the operator $\mathcal{H} = -\frac{1}{2}\nabla^2 + V(x)$, $N = 4,000$, $M = 20$, $\epsilon = 10^{-7}$, $L = 600$.

algorithm takes roughly half the numbers of iterations.

5. **Conclusion.** We see that the addition of the localization step substantially improved the performance of orbital minimization, in the sense that it significantly reduced the chances that the solution be trapped at local minima. It also improves the convergence rate for the iteration process. In addition, as we will show in future publications, it substantially improves the accuracy of orbital minimization with truncation.

From an algorithmic viewpoint, the cost of the localization step is quite reasonable. In addition, one does not have to include the localization step at each step

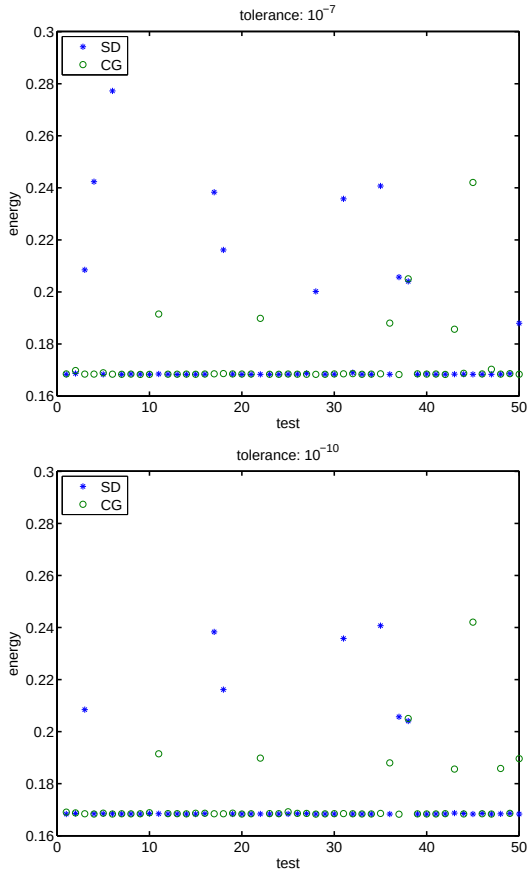

FIGURE 9. 50 random tests of the steepest descent and conjugate gradient algorithms for the second order discretization of one-dimensional Laplacian operator, $N = 500$, $M = 10$, $L = 150$, ϵ is chosen to be 10^{-7} and 10^{-10} respectively.

of the iteration. It can be used after a number of steps, or whenever the iteration shows sign of been trapped at local minima.

Many algorithmic issues remain before this becomes a reliable algorithm for electronic structure analysis. These issues include: adaptively choosing the support of the orbitals, efficient strategies for dealing with the nonlinear aspects of the problem (where H itself depends on Ψ), and extension to finite temperature. These issues will be addressed in future publications.

Acknowledgements. We are very grateful to Carlos Garcia-Cevera, Jianfeng Lu and Yulin Xuan for helpful discussions.

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*, 69(7):1077–1080, 1992.

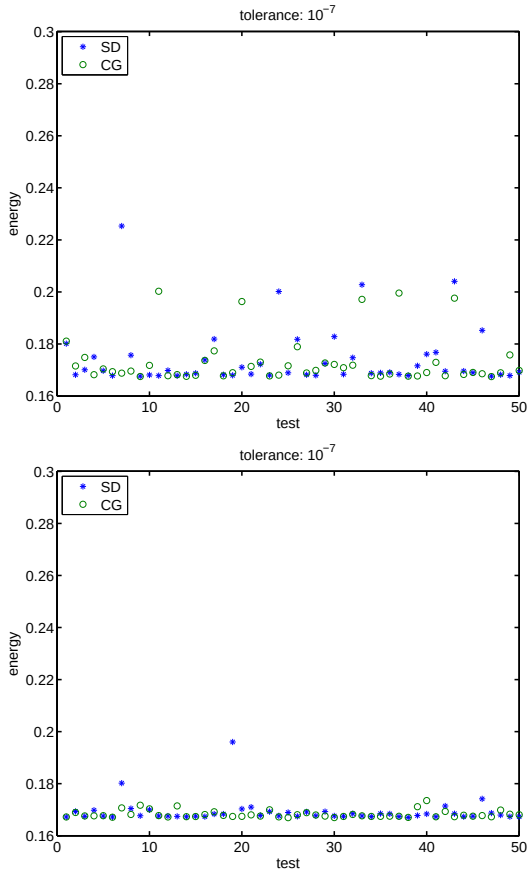

FIGURE 10. 50 random tests of the steepest descent and conjugate gradient algorithms for the second order discretization of one-dimensional Laplacian operator, $N = 500$, $M = 10$, $\epsilon = 10^{-7}$, L is set to be 250 and 350 respectively.

[2] W. E, T. Li, and J. Lu. Localized basis of eigen-subspaces and operator compression. submitted.
[3] A. Edelman, T. A. Arias, and S. T. Smith. The geometry of algorithms with orthogonalilty constrains. *SIAM Journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.
[4] G. Galli. Large-scale electronic structure calculations using linear scaling methods. *Physica Status Solidi (B)*, 217(1):231–249, 2000.
[5] G. Galli and M. Parrinello. Large scale electronic structure calculations. *Physical Review Letters*, 69(24):3547–3550, 1992.
[6] C. Garcia-Cevera and J. Lu. private communication.
[7] S. Goedecker. Linear scaling electronic structure methods. *Reviews of Modern Physics*, 71(4):1085–1123, 1999.
[8] G. H. Golub and C. van Loan. *Matrix computations*. The Johns Hopkins University Press, London, third edition, 1996.
[9] J. Kim, F. Mauri, and G. Galli. Total-energy global optimizations using nonorthogonal localized orbitals. *Physical Review B*, 52(3):1640–1648, 1995.
[10] W. Kohn. Density functional and density matrix method scaling linearly with the number of atoms. *Physical Review Letters*, 76(17):3168–3171, 1996.

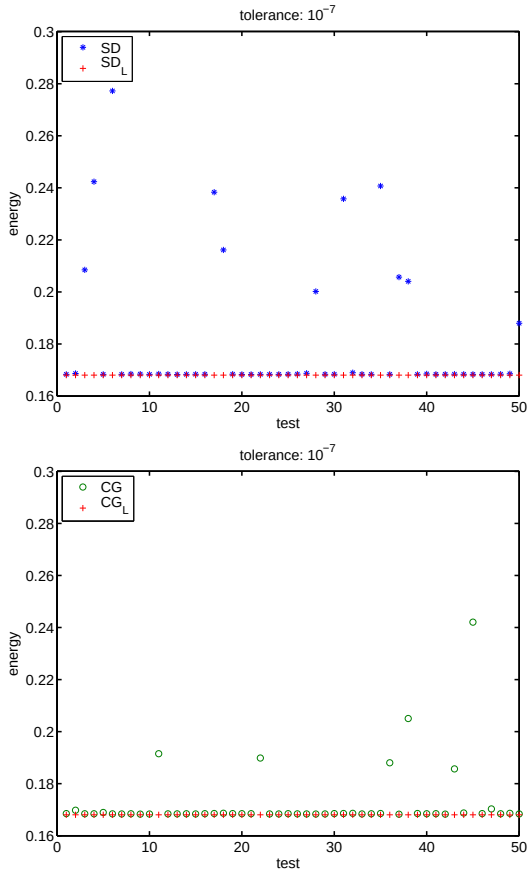

FIGURE 11. 50 random tests of the steepest descent (SD), steepest descent with localization (SD_L), conjugate gradient (CG) and conjugate gradient with localization (CG_L) algorithms for the second order discretization of one-dimensional Laplacian operator, $N = 500$, $M = 10$, $\epsilon = 10^{-7}$, $L = 150$.

[11] R. A. Lippert and M. P. Sears. Asymptotic convergence for iterative optimization in electronic structure. *Physical Review B*, 61(19):12772–12785, 2000.

[12] N. Marzari and D. Vanderbilt. Maximally localized generalized wannier functions for composite energy bands. *Physical Review B*, 56(20):12847–12865, 1997.

[13] F. Mauri and G. Galli. Electronic-structure calculations and molecular-dynamics simulations with linear system-size scaling. *Physical Review B*, 50(7):4316–4326, 1994.

[14] F. Mauri, G. Galli, and R. Car. Orbital formulation for electronic-structure calculations with linear system-size scaling. *Physical Review B*, 47(15):9973–9976, 1993.

[15] P. Ordejón, D. A. Drabold, M. P. Grumbach, and R. M. Martin. Unconstrained minimization approach for electronic computations that scales linearly with system size. *Physical Review B*, 48(19):14646–14649, 1993.

[16] P. Ordejón, D. A. Drabold, R. M. Martin, and M. P. Grumbach. Linear system-size scaling methods for electronic-structure calculations. *Physical Review B*, 51(3):1456–1476, 1995.

[17] B. G. Pfrommer, J. Demmel, and H. Simon. Unconstrained energy functionals for electronic structure calculations. *Journal of Computational Physics*, 150:287–298, 1999.

[18] E. Prodan and W. Kohn. Nearsightedness of electronic matter. *Proceedings of the National Academy of Sciences*, 102(33):11635–11638, 2005.

[19] D. Raczkowski, C. Y. Fong, P. A. Schultz, R. A. Lippert, and E. B. Stechel. Unconstrained and constrained minimization, localization, and the Grassmann manifold: Theory and application to electronic structure. *Physical Review B*, 64(15):155203, 2001.

[20] E. B. Stechel, A. R. Williams, and P. J. Feibelman. n-scaling algorithm for density-functional calculations of metals and insulators. *Physical Review B*, 49(15):10088–10101, 1994.

[21] W. Yang. Absolute-energy-minimum principles for linear-scaling electronic-structure calculations. *Physical Review B*, 56(15):9294–9297, 1997.

Received September 2006; revised February 2007.

E-mail address: wggao@fudan.edu.cn
E-mail address: weinan@math.princeton.edu