

Response to Reviewer 1

Summary Comment

The authors have addressed many of the issues raised and the manuscript is much improved. There remain some significant issues, however. If these can be well addressed, I would recommend publication.

Response:

We would like to express our sincere gratitude for your feedback and valuable suggestions on our manuscript. We have carefully considered each of your points and revised them carefully.

Comment 1.1

Regarding what is new and significant in the present work, having now included related works Refs. 33 and 34 in the discussion, it is clear that the main novelty of the present work is the treatment of both molecules and solids using the same general framework, and the basis-set independent mathematical formulation enabling the use of orthogonal as well as non-orthogonal basis sets in discretization. However, while this is sufficiently clear from the responses to reviewers, it is not sufficiently clear in the manuscript. Given the overlap of the present work with Refs. 33 and 34, what is new and significant in the present work relative to Refs. 33 and 34 should be clarified specifically. Regarding Ref. [33], the manuscript says, "In these works by ... Dai et al. [33], finite systems were treated and periodic systems were not seriously considered." What is meant by "seriously considered"? What specifically is new and significant in the present work relative to Ref. [33] in this regard? Regarding Ref. [34], the manuscript says, "We note that the recent 2023 paper targeted for metallic system in a planewave basis [34] aligns closely with our approach." What is meant by "aligns closely"? What specifically is new and significant in the present work relative to Ref. [34]? These should be clarified in the manuscript. Also, mentioning the possibility of using the proposed RCG methodology in RDMFT calculations, but not actually doing it in the present work, is not itself a sufficient distinction to warrant publication of the present work.

Response 1.1:

Thank you for pointing out the potential confusion in the manuscript. In the quoted text, we wrote: "In these works by Wen et al. [29], Zhang et al. [30], and Dai et al. [33], finite systems were treated, and periodic systems were not seriously considered."

Regarding Ref. [33], the phrase "not seriously considered" means that the work did not address periodic systems at all, and as a result, the underlying manifold is real (as discussed in the text following Eq. 2.18 in their paper). All numerical experiments in Ref. [33] were conducted on finite systems.

Regarding Ref. [34], the phrase "aligns closely" refers to the fact that their work operates on an intrinsically complex manifold, which arises from the complexity of Bloch states. In this aspect, Ref. [34] is the closest to our work. However, their approach is specifically motivated by metallic systems, with significant effort devoted to the dual-step scheme. Additionally, a key distinction is that Ref. [34] employs a plane-wave basis, whereas our current implementation uses a localized basis.

Action:

We have modified the text regarding Refs. [33] and [34] such it now reads " Ref. [33] did not address periodic systems. As a result, the underlying manifold is thus real in principle and only when the basis (e.g. planewaves) is itself complex does the manifold become complex. However, in the periodic systems, complexity is required due to the intrinsic complexity of the Bloch states. We note that a recent 2023 paper [34], which is

motivated by metallic systems and employs a plane-wave basis [34], aligns closely with our approach in this regard, despite our use of a localized basis in the current work.”

Comment 1.2

The results shown for the Riemannian Broyden-Fletcher-Goldfarb-Shanno (RBFGS) method, as an alternative to the Riemannian conjugate gradient (RCG) method, are interesting. However, there is no mathematical formulation, algorithm, or implementation information given. If the RBFGS method is to be included in the present manuscript, then a mathematical formulation, algorithm, and implementation information should be provided. Otherwise, the reader has no way to know what has actually been done, where RBFGS is the same as RCG and where it is different, and so no basis to understand the results presented.

Response 1.2: Thank you for your suggestion.

We would like to clarify that we have also implemented the Riemannian steepest descent (RSD) method. However, we chose to omit this part from the manuscript as we believe it is of limited interest compared to the other methods discussed. If you feel it would add value to the paper, we would be happy to include it in a revised version. Regarding BFGS, we have included the formulations and algorithm for RBFGS in the appendix. We have also modified the title to reflect the inclusion of the RBFGS method.

Comment 1.3

Regarding Algorithm 1 and Figure 1, because the most general case includes multiple k-points, and it is not entirely clear from the present discussion how these are incorporated into the proposed method, Algorithm 1 and Figure 1 should include multiple k-points.

Response 1.3:

To apply to the multiple k-points cases, one simply expands the dimension of the pertinent X . This is achieved by stacking K copies of X , each of the same size, to accommodate the additional k-points.

Action:

We have added description on how to include k -points.

Comment 1.4

Regarding Algorithm 1, because it is stated in terms of compact function notation, with subscripts and superscripts as well as arguments, it remains unclear to me exactly what linear algebra operations the notations correspond to. In order to clarify what the method is doing, the present compact function notation, e.g., $R_{X_k} \alpha_k d_k$, should be clarified by including the corresponding linear algebra operations performed. For example, rather than just $R_{X_k} \alpha_k d_k$ the algorithm should have $R_{X_k} \alpha_k d_k = qf(X_k + \alpha_k d_k)$ where $qf(X_k + \alpha_k d_k)$ is the Q factor of the QR decomposition of $X_k + \alpha_k d_k$. (Assuming I've interpreted the compact notation correctly in this case.) Similar explicit linear algebra representations should be given for all other such compact function notations as well, $\text{grad } f(X_{k+1})$ and $T_{\alpha_k d_k}(d_k)$ in particular. Otherwise, it's difficult to be sure exactly what linear algebra operations are taking place and in what order.

Response 1.4: Thank you for your valuable suggestions.

Your interpretation of the notation is entirely correct, and we believe that it is clear upon careful reading. However, we think the current notation should remain general, as it allows for flexibility in incorporating alternative retractions and vector transports. To ensure clarity, we have included an illustrative example in the manuscript, which will enable readers to adapt the notation to their specific needs without difficulty.

Action:

We have added after Algorithm 1, “For $\mathcal{R}_{X_k}(\alpha_k d_k) = \text{qf}(X_k + \alpha_k d_k)$, $\text{qf}(X_k + \alpha_k d_k)$ is the Q factor of the QR decomposition of $X_k + \alpha_k d_k$. For other retractions and vector transports, one needs to apply the corresponding linear algebra operations.”

Comment 1.5

Nonlocal SG15 ONCV pseudopotentials are used but at no point in the exposition did I find where and how the nonlocality is taken into account, i.e., operations with nonlocal projectors. This should be clarified.

Response 1.5:

For ONCV, the nonlocality of pseudopotential is included following the standard nonlocal projectors,

$$\hat{v}_\alpha^{NL} = \sum_{l=0}^{l_{\max}} \sum_{m=-l}^l \sum_{n=1}^{n_{\max}} |\chi_{\alpha l m n}\rangle \langle \chi_{\alpha l m n}|$$

where $|\chi_{\alpha l m n}\rangle$ are non-local projectors, with l, m, n being the angular momentum, the magnetic momentum, and the multiplicity of projectors, respectively. l_{\max} and n_{\max} are the maximal angular momentum and the maximal multiplicity of projectors for each angular momentum channel.

Action:

We have modified the pseudopotential description to read “The nonlocality of pseudopotential is included via the standard nonlocal projectors.”

Comment 1.6

In their reply, the authors state, “We believe the critical issue is not about how accurate the calculation we try to achieve.” The accuracy required is a major factor in the selection of a method of solution. One method may converge more quickly than another to low accuracy but less quickly to high accuracy (e.g., polynomial basis vs. planewave). A method that produces ill-conditioned linear systems may converge rapidly to low accuracy but encounter difficulties getting to higher accuracies. So, when comparing different methods, it is important to keep in mind the accuracies required. In the case of Kohn-Sham electronic structure calculations, that accuracy is typically on the order of chemical accuracy (1e-3 Ha/atom). Hence a method capable of achieving accuracy on this order efficiently relative to other methods will be of interest for KS calculations whether or not it achieves much higher accuracies efficiently relative to other methods. This is why it is important to target accuracies typically targeted in practice when proposing and comparing a new methodology.

Response 1.6:

We believe there may have been some misunderstanding in our earlier reply, and we would like to clarify our position. We fully agree with your comment on the importance of accuracy in the proposed methodology. What we intended to convey is that our method is capable of reaching a converged solution with the desired target accuracy, provided a sufficient number of iterations are allowed.

As you rightly pointed out, when comparing different methods, the only meaningful measure of efficiency is to ensure that the comparisons are made at the same level of accuracy. With this in mind, we have made every effort to achieve and report results with similar accuracy across all methods in our study.

Comment 1.7

The authors have added a helpful flowchart (Fig. 1) of the RCG method. However, there is no corresponding flowchart in the manuscript for the SCF method being compared to. It would be most helpful for the authors to clarify which steps are the same and which are different in the SCF method being compared to. The authors clarify this in the reply but not in the manuscript. It should be clarified in the manuscript.

Response 1.7: Thank you for your suggestion.

We initially did not include a flowchart for the SCF method, as it follows a fairly standard and well-established scheme. However, in response to your request, we have now added the flowchart to the manuscript as an appendix for clarity and completeness.

Comment 1.8

Regarding the mixing method and parameters used in the SCF calculations, there are a few concerning statements in the Response to reviewers and manuscript:

- a) "... mixing parameter `mixing_beta=0.8` (for `nspin=1`) were used." In practice the mixing parameter is tuned to decrease SCF iterations. Thus comparisons to the SCF method without tuning the mixing parameter may not be meaningful.
- b) "By default, the Kerker preconditioner is also turned on." (Manuscript p. 9) The Kerker preconditioner targets metallic and other such systems with significant fractional orbital occupations, e.g., occupation significantly less than 2 for one or more orbitals in non-spin-polarized calculations. It should not be used by default, especially for systems without significant fractional occupation. Any results using Kerker for systems without significant fractional occupation should be re-run without Kerker – and mixing parameter decreased as needed to attain efficient SCF convergence.
- c) "...the number of iterations with Kerker preconditioner on by default for molecular system is slightly larger (1-2 more iterations) than when we turned off the preconditioner" This is to be expected. The Kerker preconditioner is targeted for metallic systems and other such systems with significant fractional orbital occupation.
- d) "However, turning off the preconditioner, three cases in SCF fail to converge (CCH radical, CH₃CH₂O radical and Si₂)." In such case, the mixing parameter should generally be reduced, rather than use the Kerker preconditioner, which may be expected to negatively affect convergence for such systems. The point of all of the above is to ensure that comparisons of the proposed method to the standard SCF method are reasonably fair, and so indicative of what can be expected in actual practice.

Response 1.8: Thank you for your insightful questions and suggestions.

Ideally, to ensure a fair comparison between different algorithms, each method should be fine-tuned with an optimal set of parameters. However, in practice, this becomes challenging as the optimal parameters can vary significantly depending on the system under study. For this reason, we followed established practices from previous works and adhered to a fixed set of parameters for consistency.

We are fully aware of the utility of the Kerker preconditioner, which is particularly effective for metallic systems and others with significant fractional occupations. In our study, we conducted

tests both with and without the Kerker preconditioner and observed that the impact on the number of iterations was minimal (up to 2 iterations) for both molecules and solids. Its default activation in ABACUS must have been done for a reason.

We agree that tuning the mixing parameter could help address cases where convergence fails. Additionally, while it would be ideal to use the same preconditioner for all methods, this has not yet been implemented in our current work. Nevertheless, we believe our results provide meaningful and informative insights, even if they are not perfect in an absolute sense.

Comment 1.9

Regarding the molecular calculations, the authors use periodic boundary conditions (BCs) rather than Dirichlet BCs as employed in practice. This may lead to undesirable long range interactions and associated difficulties in SCF convergence. It will also generally require a significantly larger domain ("box") size to reduce associated finite-size errors to typical $1e-3$ Ha/atom targets. To reduce the need for excessively large box size (and associated computational cost and SCF convergence issues), molecular calculations typically employ vanishing Dirichlet BCs for the wavefunctions and charge density, and non-vanishing Dirichlet BCs for the potential at the domain boundary (with potential on the boundary computed via multipole expansion). Non-vanishing BCs for the potential are necessary since it decays less rapidly than the wavefunctions and density (e.g., as $1/r^2$ for a neutral molecule). In order to much better reflect actual practice, it would strengthen the manuscript considerably to perform at least a small subset of molecular calculations with Dirichlet BCs and correspondingly smaller boxes (e.g., just large enough to reduce associated energy errors to $1e-3$ Ha/atom) in order to better show how the RCG method can be expected to compare to the standard SCF method in molecular calculations in practice.

Response 1.9: Thank you for your valuable comments.

We would like to clarify that the current implementation of the code does not include vanishing Dirichlet boundary conditions (BCs). Implementing such boundary conditions falls outside the scope of our present work. We appreciate your understanding and welcome further discussions on this topic for future developments.

Comment 1.10

Regarding timings, the breakdown shown in the reply is most helpful. A breakdown like this should be shown in the manuscript for both SCF and RCG methods for the same problem and convergence/stopping criteria. This would make clear not only which operations are the same and which are different in the proposed RCG method relative to the standard SCF method, but also clarify the relative cost of those operations. I noticed one issue in the breakdown, however. The time for diagonalization and density construction in the SCF method is shown to be 0.78 s per call for the $\text{CH}_2=\text{CH}-\text{CH}=\text{CH}_2$ system. Diagonalization for such small systems should be much faster. Indeed, diagonalization of a 500x500 matrix takes less than 0.1 seconds on a single (Intel Cascade Lake Xeon) core using LAPACK. Since 500 orbitals is many more than required for the $\text{CH}_2=\text{CH}-\text{CH}=\text{CH}_2$ system, it seems something is significantly inefficient in the diagonalization step of the SCF implementation. And since the diagonalization step is the main difference between the SCF and RCG methods, this should be clarified. Also, the potential calculation (`updatepot()`) seems rather long for such a small Poisson (linear system) solution. Is this step the same in the RCG implementation? If not, why not (since both must construct the KS potential)?

Response 1.10:

We have followed your suggestions and made a short note of the relative cost of each operation. Regarding the time of diagonalization, in the step of `hamilt2density()` additional operations in forming the density on grid are performed and the diagonalization takes less time. Also in the

step of `updatepot()`, in addition to the Poisson solver, each potential is calculated and the relevant Hamiltonian is updated in both real space and reciprocal space. This step is shared by both RCG and SCF methods.

To provide further clarification, the majority of the computational time in the RCG method is spent during the line-search step. This involves evaluating the objective function and its gradient, which rely heavily on numerical integrals and represent the most resource-intensive part of the process.

Comment 1.11

Regarding stopping criteria, for a meaningful comparison of SCF and RCG methods, the same stopping criteria must be used. If one method is required to converge to a greater degree than the other, then it will take more iterations and time to do so, thus changing the results of the comparison. As currently implemented, the SCF method is using a charge density convergence criterion while the RCG method is using an energy convergence criterion. Simplest and most directly meaningful would be for both SCF and RCG methods to use the same energy convergence criterion with tolerance corresponding to that typical in practical calculations. For example, stop when the difference of 3 successive total energies is less than $1e-5$ Ha/atom. (3 rather than 2 successive for additional robustness in the absence of a charge density convergence criterion to verify.)

Response 1.11:

We agree that using the same energy convergence criterion is a best practice for fair and accurate comparisons. Given the energy accuracy achieved by both the RCG and SCF methods in our study, we believe the comparison remains justified. As mentioned in our previous reply, the energy difference criterion was only implemented in a later version of the code.

We would like to note that the latest version of ABACUS has undergone significant code restructuring. We are currently in the process of adapting our implementation to align with these changes and integrating our modifications into the main branch. As part of this effort, we are also working to minimize redundant function calls to improve efficiency.

We greatly appreciate your suggestion regarding the use of three successive energies for convergence assessment, and we are happy to incorporate this into our implementation moving forward.

Comment 1.12

Finally, it would strengthen the manuscript considerably if the authors could show and discuss in detail at least one case where the RCG method is significantly faster than the standard SCF method, ensuring computational resources are the same, SCF parameters and preconditioners are reasonably chosen, stopping criteria are the same, and accuracy targeted is $1e-3$ Ha/atom in the final total energy (absolute error with respect to all computational parameters, e.g., as compared to a well converged reference result from literature or standard code). If such a case cannot be found, then detailed discussion of any other case where RCG does reasonably well would be informative as well. Discussion of the timing breakdown for all major steps of the RCG and SCF methods would then make much clearer, quantitatively, what the relative strengths and weaknesses of the RCG method are, and clarify the prospects for further improvements and practical impact going forward.

Response 1.12: We sincerely appreciate your valuable suggestions.

Previous studies have demonstrated that the RCG method offers significant advantages over the SCF method, particularly as the system size increases. This advantage was further confirmed through our modification of the MATLAB-based code, KSSOLV, using the RCG approach. We believe that highlighting a single specific case may not be entirely convincing. We would project that the RCG method, when equipped with a more intelligent line search scheme and an improved

preconditioner, can significantly reduce the number of iterations required for convergence.

Moreover, in scenarios where SCF methods struggle to converge, the RCG method in principle should be a robust alternative, comparable to established minimization methods such as the band-by-band CG and the all-band CG. Based on the promising results from this pilot implementation, we are confident that the RCG method holds substantial potential and will demonstrate its value in broader applications.

Response to Reviewer 2

Summary Comment

The manuscript reports reports from testing an electronic minimizer using the Stiefel manifold differential geometry. I like the tests and comparisons to SCF minimizers the study reports. However, the manuscript leaves out details about their minimizer(s) that would make it hard to the reader to directly implement there method. While some of these details can be found in the work of Edelman et al., and potentially algorithms implemented the nwchem pw, many of the details are not provided. In particular, they use a bfgs quasi-newton method which is not very well explained and details like the number of histories used and the parallel transport for the iterative hessians are not provided. That being said, I think the goal of the manuscript is how the performance of the bfgs stiefel, cg stiefel, and more standard scf methods compare to each other, but it would be nice to be able to provide access to codes so that readers can compare themselves. Also esoteric differential geometry math is probably not a priority for readers of this journal, which make writing this type of manuscript difficult. Since this is a journal focused on presenting algorithms this should be possible.

Response: We sincerely thank you for your positive assessment and constructive suggestions.

In response to your feedback, we have added more detailed explanations regarding the minimizer. A new reference to the GitHub is added. Additionally, as suggested by Reviewer 1, we have included an appendix that provides further insights into the RBFSG method. We have also modified the title to reflect the inclusion of the RBFSG method.

To facilitate broader testing and collaboration, we have made the current version of the code available on GitHub at the following link: https://github.com/kluophysics/directmin_abacus.git. We would like to note that this code is still under active development and will eventually be merged into the main branch. We encourage users to explore and test the code, and we welcome any feedback to further refine its functionality.

We have carefully considered your suggestion regarding the omission of the esoteric differential geometry content in the subsection on the complex Stiefel manifold. While we acknowledge that moving the background material on the Stiefel manifold to the appendix could reduce the reading burden—particularly given the technical nature of retraction and vector transport—we are concerned that doing so might compromise the coherence and completeness of these concepts. After thorough deliberation, we have decided to retain this subsection in its current form to ensure that readers who wish to explore this topic in depth have access to a comprehensive and self-contained presentation.

We have also fixed many typos.